arXiv:1603.05725v1 [math.GR] 17 Mar 2016

ACYLINDRICAL HYPERBOLICITY OF CUBICAL SMALL-CANCELLATION GROUPS

GOULNARA N. ARZHANTSEVA AND MARK F. HAGEN

Abstract. We provide an analogue of Strebel's classification of geodesic triangles in classical $C'(\frac{1}{6})$ groups for groups given by Wise's cubical presentations satisfying sufficiently strong metric cubical small-cancellation conditions. Using our classification, we prove that, except in specific degenerate cases, such groups are acylindrically hyperbolic.

Contents

Introduction 1
1. Acylindrical hyperbolicity and WPD elements 6
2. Triangles in cubical small-cancellation groups 7
3. Detecting WPD elements using the classification of triangles 16
4. The hyperbolic space $\mathcal{H}$ and the coarse projection 19
5. Proof of Theorem B 20
References 25

Introduction

A cubical presentation of a group is a natural high-dimensional generalization of both a "classical" and a "graphical" presentation of a group in terms of generators and relators. The theory of cubical presentations, and especially the cubical small-cancellation theory developed by Wise [Wis, Section 3], has begun to play a significant role in geometric group theory following spectacular solutions of the virtual Haken conjecture by Agol and of Baumslag's conjecture on one-relator groups with torsion by Wise. Topologically, a *classical presentation* of a group G consists of a wedge X of circles and a collection of combinatorial immersions $Y_i \to X$ of circles so that the *presentation complex* X^* formed from X by coning off the various Y_i satisfies $\pi_1 X^* \cong G$. The 1–skeleton $\mathrm{Cay}(X^*)$ of the universal cover $\widetilde{X}^*$ of X^* is a Cayley graph of G with respect to the generating set implicit in the choice of X. A *graphical presentation* is a natural generalization of this: X is allowed to be an arbitrary graph, and each $Y_i \to X$ becomes an immersion of graphs.

In [Wis], it is observed that allowing even more flexibility in the choice of X can lead to more tractable "presentations" for a given group. This leads to the notion of a *cubical presentation*: X is now a nonpositively-curved cube complex and each Y_i is a connected nonpositively-curved cube complex equipped with a local isometry $Y_i \to X$. The presentation complex X^* is defined analogously, and there is a *generalized Cayley graph* $\mathrm{Cay}(X^*)$ which is the cubical part of the universal cover of X^*, i.e. the cover of X corresponding to $\ker(\pi_1 X \to \pi_1 X^*)$. The analogy with classical presentations is clear: the cube complex X is a kind of "high-dimensional generating

2010 *Mathematics Subject Classification.* 20F06, 20F65, 20F67.
Key words and phrases. Cubical small-cancellation, acylindrically hyperbolic group.
G.N. Arzhantseva was partially supported by her ERC grant ANALYTIC no. 259527.
M.F. Hagen was supported by the EPSRC grant of Henry Wilton.

set", the CAT(0) cube complex $\widetilde{X}$ is the "high-dimensional tree" taking the place of the free group on the generating set in the classical case, and $\mathrm{Cay}(X^*)$ corresponds to a Cayley graph. The advantage is that some of the complication of the group has been "stored" in the organized, tractable cubical structure of X, leaving us to worry only about the interactions between the various Y_i, and between each Y_i and the cubical structure of X.

The latter problem is handled by cubical small-cancellation theory, in which "generalized overlaps" between the various Y_i (i.e. shadows of Y_i on Y_j, as propagated through the intervening cubes) are small in the appropriate metric sense. In this setting, there are powerful tools – specifically, the *ladder theorem* and the *cubical Greendlinger lemma* (see Section 2) – that allow one to extract considerable geometric and algebraic information about a group from a small-cancellation cubical presentation. The small-cancellation conditions of interest in this paper are the *cubical* $C'(\alpha)$ *conditions*, for $\alpha > 0$. These say that $|P| < \alpha\|Y_i\|$ for all P, i, where $\|Y_i\|$ denotes the length of a shortest essential closed combinatorial path in Y_i and $|P|$ is the length of the geodesic *piece* P. A piece is a path in the "generalized overlap" between distinct elevations to $\widetilde{X}$ of the various Y_i, or between such elevations and hyperplane-carriers in $\widetilde{X}$.

One advantage of passing to cubical presentations is that many groups that do not admit classical presentations satisfying strong small-cancellation conditions nonetheless admit cubical presentations with these properties. For example, if G is the fundamental group of a nonpositively-curved cube complex X, then G admits a cubical presentation with no relators, and therefore satisfies *arbitrarily* strong cubical small-cancellation conditions; on the other hand, G does not in general satisfy strong classical small-cancellation conditions, as can be seen by considering, for instance, right-angled Artin groups. Later in this introduction, we list more examples of cubical small-cancellation groups.

Our first result is geometric. We classify geodesic triangles in the generalized Cayley graph of a cubical small-cancellation group in terms of the disc diagrams that they bound in the presentation complex. This is a cubical analogue of Strebel's classification of geodesic triangles in $C'(\frac{1}{6})$ groups (Theorem 43 of [Str90]), which says that any geodesic triangle bounds a disc diagram of one of a small number of very specific combinatorial types:

Theorem A (Classification of triangles)**.** *There exists $\alpha > 0$ so that the following holds. Let X be a nonpositively-curved cube complex, let $\mathcal{I}$ be a (possibly infinite) index set and let $\{Y_i \to X : i \in \mathcal{I}\}$ be a set of local isometries. Suppose that the cubical presentation $\langle X \mid \{Y_i\}_{i\in\mathcal{I}}\rangle$ satisfies the cubical $C'(\alpha)$ condition. Let X^* be the presentation complex and x, y, z be 0–cells of the universal cover $\widetilde{X}^*$. Then there exists a geodesic triangle Δ in $\widetilde{X}^*$, with corners x, y, z, so that Δ is the boundary path of a disc diagram $D \to X^*$ of one of* 9 *types; in particular, D is the union of* 3 padded ladders. *Moreover, any other geodesic triangle with corners x, y, z is square-homotopic to Δ.*

The precise (longer) statement is Theorem 2.16, which explains exactly what the "9 types" of disc diagram are; a *padded ladder* is a disc diagram of the type in Figure 2. The α required in our proof is $\frac{1}{144}$. Conceptually, our theorem says that any geodesic triangle bounds a disc diagram which is square-homotopic (fixing corners) to a disc diagram which is a "thickened tripod".

Theorem A can be used to recover existing results of a similar character by imposing various conditions on X or $\mathcal{I}$.

- If $\mathcal{I} = \emptyset$, then Theorem A says that any three vertices in a CAT(0) cube complex determine a geodesic tripod, which is a consequence of the fact that CAT(0) cube complexes are exactly the simply-connected cube complexes whose 1–skeleta are median graphs [Che00].
- If X is a wedge of circles and each Y_i is an immersed circle, then $\langle X \mid \{Y_i\}\rangle$ is a classical $C'(\frac{1}{144})$ presentation, and the original Strebel classification for classical $C'(\frac{1}{6})$ groups applies, and follows from Theorem A.

- If $\dim X = 1$ (i.e. X is a graph) and each $Y_i \to X$ is an immersion of graphs, then we have graphical presentations, which are the setting for graphical small-cancellation theory. In this setting, there is a classification of triangles that holds under weaker small-cancellation conditions than are required in the cubical setting. Indeed, the classification of triangles is completely combinatorial, and Strebel's proof actually applies in the setting of the $(3,7)$–*diagrams* used by Gruber-Sisto in their proof of acylindrical hyperbolicity for graphical small-cancellation groups [GS14]; this combinatorial observation was made by Gruber [Gru15, Remark 3.11]. While the result about $(3,7)$–diagrams suffices for graphical small-cancellation groups, one cannot extend it directly to disk diagrams over cubical presentations since the presence of squares means that such diagrams need not satisfy the $(3,7)$ condition.

We foresee several applications of Theorem A to the thorough investigation of cubical small-cancellation groups (cf. in classical small-cancellation theory, the classification of triangles is crucial in calculating the conformal dimension of the boundary [Mac], in proving growth tightness [Sam], the rapid decay property [AD12], SQ-universality [Gru], etc.).

In this paper, we focus on acylindrical hyperbolicity, inspired by the corresponding result for graphical small-cancellation groups [GS14]. A group G is *acylindrically hyperbolic* if it admits a nonelementary acylindrical action on a hyperbolic space (acylindricity generalizes *uniform* properness in a natural way). The notion of acylindrical hyperbolicity, due to Osin [Osi15a], unifies several generalizations of relative hyperbolicity [BF02, DGO11, Ham08, Sis] and provides a class of groups with many strong properties: if G is acylindrically hyperbolic, then G is SQ-universal, contains normal free subgroups, and is C^*–simple if and only if it has no finite normal subgroup [DGO11]; G contains Morse elements and thus all asymptotic cones of G contain cut-points [Sis14]; the bounded cohomology of G has infinite dimension in dimensions 2 [HO13] and 3 [FPS15]; every commensurating endomorphism of G is an inner automorphism [AMS], etc.

The class of acylindrically hyperbolic groups is now known to be vast, see e.g. [Bow08, Osi15a, DGO11, MO15, Osi15b, BF08, GS14, BHS14, PS15]. Our second result adds Wise's cubical small-cancellation groups to this notable list:

Theorem B (Acylindrical hyperbolicity from cubical small-cancellation)**.** *Let X be a compact nonpositively-curved cube complex and let $\langle X \mid \{Y_i\}_{i\in\mathcal{I}}\rangle$ be a (possibly infinite) $C'(\frac{1}{144})$ cubical presentation with each Y_i compact. Let X^* denote the presentation complex. Then one of the following holds:*

(1) $\pi_1 X^*$ is finite;
(2) $\pi_1 X^*$ is two-ended;
(3) each Y_i is contractible, $\pi_1 X^* = \pi_1 X$, and the universal cover $\widetilde{X}$ of X contains a convex $\pi_1 X$–invariant subcomplex splitting as the product of unbounded cube complexes;
(4) $\pi_1 X^*$ is acylindrically hyperbolic.

In the case where X is a compact cube complex and $\mathcal{I} = \emptyset$ (i.e. in the *purely cubical* case), the failure of $\pi_1 X$ to be hyperbolic corresponds to large complete bipartite subgraphs of the *contact graph* – the intersection graph of the hyperplane-carriers – of $\widetilde{X}$, as shown independently in [Hag14] and [Che00]; the space $\mathcal{H}$ obtained from $\widetilde{X}$ by coning off the hyperplane-carriers is always hyperbolic [Hag14]. In [BHS14], it is shown that under natural extra hypotheses, the action of $\pi_1 X$ on $\mathcal{H}$ is acylindrical, while in any case any $g \in \pi_1 X$ acting loxodromically on $\mathcal{H}$ actually acts as a WPD element in the sense of [BF02] (see also Definition 1.2 below); together with results in [Hag13] characterizing the loxodromic isometries of $\mathcal{H}$, and a result of Osin connecting WPD elements to acylindricity [Osi15a], this implies the virtually cyclic/product/acylindrically hyperbolic trichotomy of Theorem B in the case where $\mathcal{I} = \emptyset$. This trichotomy (in the purely cubical case) also follows from the Caprace-Sageev rank-rigidity theorem [CS11] and general results about groups acting on CAT(0) spaces and containing rank-one elements [Osi15a, Sis],

although the first argument seems more amenable to generalization along the lines of Theorem B (since that is what we do in the present paper).

The comparison with the acylindrical hyperbolicity result of Gruber-Sisto, for graphical small-cancellation groups, is interesting; our results about cubical small-cancellation groups do not follow from corresponding results about graphical small-cancellation presentations, since the latter viewpoint does not account for high-dimensional cubes. On the other hand, restricting Theorems A and B to the case where $\dim X = 1$ and each Y_i is a graph, one does not reprove the results of [GS14] or [Str90] in full generality, since the cubical $C'(\frac{1}{144})$ condition is a priori more restrictive than the small-cancellation conditions needed in the classical and graphical cases (which are the classical $C'(\frac{1}{6})$ and the graphical $Gr(7)$ conditions, respectively).

Theorem B is proved roughly as follows. First, we create a hyperbolic $\pi_1 X^*$–space $\mathcal{H}$ by coning off each hyperplane-carrier $\mathcal{N}(H)$ and each relator Y_i in the generalized Cayley graph $\mathrm{Cay}(X^*)$. This procedure is a common generalization of the constructions used in the purely cubical case (where the space obtained from coning off the hyperplane-carriers is quasi-isometric to the contact graph) [Hag14] and in the graphical case (where the space is obtained from coning off each relator graph Y_i by attaching the complete graph on its vertices) [GS14]. Next, we apply Theorem A to show that if $g \in \pi_1 X^*$ acts loxodromically on $\mathcal{H}$, then g is a WPD isometry of $\mathcal{H}$, and therefore Osin's theorem tells us that $\pi_1 X^*$ is acylindrically hyperbolic or virtually cyclic. It remains to find loxodromic isometries of $\mathcal{H}$. This is done in Section 5 over a series of lemmas. Roughly, we show that if $\tilde{g} \in \pi_1 X$ acts loxodromically on the contact graph of $\widetilde{X}$ (and whether there are such $\tilde{g}$ is well-understood [Hag13]), and $\tilde{g}$ survives in the quotient $\pi_1 X \to \pi_1 X^*$, then the image g of $\tilde{g}$ has the desired properties. It is then a matter of using the small-cancellation assumptions to show that either no such $\tilde{g}$ was available (so, roughly, $\widetilde{X}$ was a product), or some $Y_i \to X$ is π_1–surjective, or there is some WPD element $g \in \pi_1 X^*$.

Remark 1 (Compactness)**.** It is in the parts of the proof dealing with product decompositions that compactness of X plays a role, since we need to invoke tools from [CS11] that require either a proper cocompact action on a cube complex, or some more exotic hypotheses about the boundary. Accordingly, one could modify the hypotheses about X and obtain a similar conclusion to that of Theorem B in various ways, mainly by weakening the statement about product decompositions. The natural conclusion if conclusions $(2)-(4)$ of the theorem fail is that there is some $Y_i \to X$ so that $\widetilde{Y}_i \subseteq \widetilde{X}$ is a $\pi_1 X$–invariant subcomplex; to get from here to conclusion (1) is the unique place where compactness of Y_i plays a role.

Remark 2 (No proof by cubulation)**.** Cubical small-cancellation theory is partly motivated by the fact that groups satisfying strong classical small-cancellation conditions act nicely on CAT(0) cube complexes [Wis04, Theorem 1.2]. This generalizes in various ways to cubical presentations: if $\langle X \mid \{Y_i\}\rangle$ satisfies the *generalized* $B(6)$ condition, one can often cubulate the corresponding group; see, for instance, [Wis, Theorem 5.50].

It is tempting to try to prove Theorem B using this approach, together with the above-mentioned results about acylindrical hyperbolicity of groups acting on cube complexes. However, there are various problems with this approach. First, the cubulations one can obtain via the generalized $B(6)$ condition may not be nice enough to invoke the results from [CS11] needed to find rank-one elements in the absence of a product decomposition. Second, the generalized $B(6)$ condition requires each Y_i to have a wallspace structure, compatible with the local isometry $Y_i \to X$, generalizing the wallspace structure on a circle in which each wall is a pair of antipodal points. (Compare with the *lacunary walling* condition on graphical presentations from [AO14].)

At the same time, no cubical $C'(\alpha)$ small-cancellation condition implies the generalized $B(6)$ condition, and indeed there are groups that are covered by Theorem B but which do not admit an action on a CAT(0) cube complex with no global fixed point. This can already be seen in the 1–dimensional case: Proposition 7.1 of [OW07] yields, for any $\alpha > 0$, a graphical presentation

$\langle X \mid Y\rangle$, where X is a graph and $Y \to X$ an immersed graph, satisfying the graphical (hence 1–dimensional cubical) $C'(\alpha)$ condition, with the additional property that the group thus presented has Kazhdan's property (T), and thus cannot act fixed-point-freely on a CAT(0) cube complex [NR98].

Remark 3 (Acylindrical action on a quasi-tree)**.** Combining Theorem B with a very recent result of Balasubramanya [Bal16] shows that any cubical small-cancellation group G covered by Theorem B either satisfies one of the conclusions $(1) - (3)$ or acts acylindrically and non-elementarily on a quasi-tree.

Examples of cubical small-cancellation groups. We list here fundamental examples and recipes for building examples, of cubical small-cancellation groups to which Theorem A and Theorem B apply. It is inspired by the list in Section 3.p of [Wis]:

(1) Classical $C'(\frac{1}{144})$ small-cancellation groups: if F is a free group and $\{\langle z_i\rangle\}_{i\in\mathcal{I}}$ a malnormal collection of cyclic subgroups, then there are $n_i' \geqslant 1$ so that the group $F/\langle\langle\{z_i^{n_i}\}_{i\in\mathcal{I}}\rangle\rangle$ satisfies the appropriate small-cancellation condition whenever $n_i \geqslant n_i'$ for all i.
(2) Graphical $C'(\frac{1}{144})$ small-cancellation groups.
(3) Classical/RAAG hybrid: let X be the Salvetti complex of a right-angled Artin group A, with presentation graph $\mathbb{G}$, and let $\{g_i\}_{i\in\mathcal{I}}$ be a collection of independent elements, none of which is supported on a proper join in $\mathbb{G}$ (i.e. each g_i is a rank-one isometry of $\widetilde{X}$). More generally, choose $\{\langle g_i\rangle\}_{i\in\mathcal{I}}$ to be a malnormal collection of cyclic subgroups, each of which has a convex cocompact core $\widetilde{Y}_i$ in $\widetilde{X}$. Then for each i there exists $n_i > 0$ so that, letting $Y_i = \langle g_i^{n_i}\rangle\backslash\widetilde{Y}_i$, the cubical presentation $\langle X \mid \{Y_i\}\rangle$ is a $C'(\frac{1}{144})$ presentation.
(4) More generally, let X be a *special* cube complex, in the sense of [HW08]. Let $\{Y_i \to X\}$ be a collection of local isometries of cube complexes so that the resulting cubical presentation $\langle X \mid \{Y_i\}\rangle$ satisfies the cubical $C'(\alpha)$ condition for some $\alpha > 0$. Each Y_i has residually finite fundamental group since it is special, and thus, for any $n \in \mathbb{N}$, there is a finite cover $\widehat{Y}_i \to Y_i$ with $\|\widehat{Y}_i\| \geqslant n\|Y_i\|$; thus the related cubical presentation $\langle X \mid \{\widehat{Y}_i\}\rangle$ satisfies the cubical $C'(\frac{\alpha}{n})$ condition. See [Wis, Section 3.q].
(5) Given letters x, y and $m \geqslant 1$, let $(x, y)^m$ denote the first half of the word $(xy)^m$. Consider the Artin group

$$A = \langle a_1, a_2, \ldots, a_n \mid (a_i, a_j)^{m_{ij}} = (a_j, a_i)^{m_{ij}} \text{ whenever } i \neq j\rangle.$$

(Note that we follow the convention of letting $m_{ij} = \infty$ to indicate that there is no relation between a_i, a_j.) Let $\widehat{A} = \langle a_1, a_2, \ldots, a_n \mid [a_i, a_j] \text{ whenever } m_{ij} = 2\rangle$ be the underlying right-angled Artin group, X be its Salvetti complex, and Γ be its presentation graph (a graph with a vertex for each a_i and with an edge from a_i to a_j when $m_{ij} = 2$). That is, A is a quotient of $\widehat{A}$ obtained by adding the relations $(a_i, a_j)^{m_{ij}} = (a_j, a_i)^{m_{ij}}$ when $m_{ij} \neq 2$. Observe that each element $g_{ij} = (a_i, a_j)^{m_{ij}}(a_j, a_i)^{-m_{ij}}$ of $\widehat{A}$, with $m_{ij} \neq 2$, is a rank-one isometry of $\widetilde{X}$, since, if it is supported in a join in the presentation graph Γ of $\widehat{A}$, then it is supported in a factor of that join. Hence there is a convex subcomplex $\widetilde{Y}_{ij}$ of $\widetilde{X}$ that is cocompactly stabilized by $\langle g_{ij}\rangle$, and which is just the convex hull of a combinatorial g_{ij}–axis. Let Y_{ij} be the quotient of $\widetilde{Y}_{ij}$ by the $\langle g_{ij}\rangle$–action, so that $\langle X \mid Y_{ij} \text{ whenever } 2 < m_{ij} < \infty\rangle$ is a cubical presentation for the Artin group A. Clearly, Y_{ij} has systole $2m_{ij}$, so in order to impose a condition on m_{ij} ensuring that this presentation satisfies the cubical $C'(\frac{1}{144})$ condition, we need only investigate the cone-pieces and wall-pieces.

If $\tilde{P}$ is a cone-piece between $\widetilde{Y}_{ij}$ and $\widetilde{Y}_{k\ell}$, then $|\tilde{P}| = 1$. On the other hand, if a_i, a_j lie in the link of some a_k in Γ, then any geodesic in $\widetilde{Y}_{ij}$ is a wall-piece, but otherwise wall-pieces have length $\leqslant 1$. Hence suppose that A satisfies the following:

- for all $i \neq j$, either $m_{ij} = 2$ or $m_{ij} = \infty$ or $m_{ij} > 72$;
- for all $i \neq j$ such that there exists k with $m_{ik} = m_{jk} = 2$, we have either $m_{ij} = 2$ or $m_{ij} = \infty$.

Then the above cubical presentation for A is $C'(\frac{1}{144})$ and our acylindricity theorem applies to A. There is a related recipe in Section 20 of [Wis] for building $C(6)$ cubical presentations of Artin groups (cf. [AS83]) but it is harder to see when these are $C'(\frac{1}{144})$.

In a similar manner, up to some limited worries about torsion, one should be able to produce many examples of Coxeter groups that virtually admit cubical $C'(\frac{1}{144})$ presentations; the base cube complex should be the Davis complex of the underlying right-angled Coxeter group.

Outline of the paper. In Section 1, we recall the definitions of acylindrical hyperbolicity and WPD elements, and a result of Osin connecting them. Section 2 begins with a discussion of cubical presentations, disc diagrams, and the features of cubical small-cancellation theory needed in the proof of the classification of triangles, Theorem A, which also occurs in this section. In Section 3, we give a list of conditions on a cubical small-cancellation group G acting on a hyperbolic space $\mathcal{H}$ sufficient to ensure that G contains an element acting on g as a WPD element. In Section 4, we produce such a space $\mathcal{H}$, formed from a generalized Cayley graph $\mathrm{Cay}(X^*)$ by coning off both the relators and the hyperplane-carriers. Finally, in Section 5, we explain when one can find elements of G acting loxodromically on $\mathcal{H}$; in concert with the results of the preceding two sections, this completes the proof of Theorem B.

We assume basic knowledge of CAT(0) and nonpositively-curved cube complexes and cubical presentations; familiarity with the background material in [Wis12] is sufficient.

Acknowledgments. MFH thanks the Fakultät für Mathematik of Universität Wien for hospitality during a visit in which much of the work on this project was completed, and is also grateful to Dani Wise for numerous discussions about cubical small-cancellation theory over the years. Both authors thank Kasia Jankiewicz for answering a question about short inner paths.

1. Acylindrical hyperbolicity and WPD elements

The notion of an *acylindrically hyperbolic group* was defined in [Osi15a], as follows:

Definition 1.1 (Acylindrical action, acylindrical hyperbolicity)**.** Let (X, d) be a metric space and let G act on X by isometries. Then the action is *acylindrical* if for each $\epsilon \geqslant 0$, there exists $R \geqslant 0$ and $N \in \mathbb{N}$ so that for all $x, y \in X$ for which $\mathsf{d}(x, y) \geqslant R$, we have

$$|\{g \in G : \mathsf{d}(x, gx) \leqslant \epsilon, \mathsf{d}(y, gy) \leqslant \epsilon\}| \leqslant N.$$

Let X be Gromov-hyperbolic and let G act by isometries on X. The action of G is *elementary* if the limit set of G on ∂X has at most two points. If G acts non-elementarily and acylindrically on a hyperbolic space, then G is said to be *acylindrically hyperbolic*.

When G is a cubical small-cancellation group, we will construct an explicit action of G on a hyperbolic space $\mathcal{H}$, but this will not necessarily be the action that witnesses acylindrical hyperbolicity. Instead, the action will be nice enough that G contains a *WPD isometry of* $\mathcal{H}$, which will in turn yield an acylindrical action on some other hyperbolic space. More precisely:

Definition 1.2 (WPD element [BF02])**.** Let G act by isometries on the space X. Then $h \in G$ is a *WPD element* if for each $\epsilon > 0$ and each $x \in X$, there exists $M \in \mathbb{N}$ so that

$$|\{g \in G : \mathsf{d}(x, gx) \leqslant \epsilon, \mathsf{d}(h^M x, gh^M x) \leqslant \epsilon\}| < \infty.$$

In [Osi15a], Osin showed that if G is not virtually cyclic and acts on a hyperbolic space $\mathcal{H}$, and some $g \in G$ acts on $\mathcal{H}$ as a loxodromic WPD element, then G is acylindrically hyperbolic. This fact is instrumental in the proof of Theorem 4.3 below, and thus in the proof of Theorem B.

2. Triangles in cubical small-cancellation groups

In this section, X denotes a connected nonpositively-curved cube complex with universal cover $\widetilde{X}$. When doing geometry in $\widetilde{X}$, we never use the CAT(0) metric and instead only use the usual graph metric on $\widetilde{X}^{(1)}$ in which each 1–cube has length 1 and a combinatorial path is geodesic if and only if it contains at most one edge intersecting each hyperplane of $\widetilde{X}$. We fix a (possibly infinite) index set $\mathcal{I}$, and for each $i \in \mathcal{I}$, let $Y_i \to X$ be a local isometry of connected cube complexes. Each Y_i is necessarily nonpositively-curved. Following [Wis], the associated *cubical presentation* is $\langle X \mid \{Y_i\}_{i\in\mathcal{I}}\rangle$ and the corresponding *cubical presentation complex* X^* is formed as follows. For each $i \in \mathcal{I}$, let $C(Y_i)$ be the *relator on* Y_i, i.e. the space formed from $Y_i \times [0,1]$ by collapsing $Y_i \times \{1\}$ to a point. This space has an obvious cell-structure so that $Y_i \xrightarrow{\sim} Y_i \times \{0\} \hookrightarrow C(Y_i)$ is a combinatorial embedding. For each $i \in \mathcal{I}$, we attach $C(Y_i)$ to X along $Y_i \times \{0\}$ using the above local isometry. The resulting complex is X^*. The group of our interest is defined by $G = \pi_1 X^*$. We say that $\langle X \mid \{Y_i\}_{i\in\mathcal{I}}\rangle$ is a *cubical presentation for* G.

The universal cover $\widetilde{X}^*$ of X^* is a nonpositively-curved cube complex with cones attached. Let $\mathrm{Cay}(X^*)$ be the part of $\widetilde{X}^*$ consisting only of cubes (i.e. the complement of the open cones). This is the *generalized Cayley graph* of G with the given cubical presentation. Note that there are covering maps $\widetilde{X} \to \mathrm{Cay}(X^*) \to X$; the generalized Cayley graph is the nonpositively-curved cube complex obtained by taking the cover of X corresponding to the kernel of $\pi_1 X \to \pi_1 X^*$.

Remark 2.1. (Classical and graphical presentations) When X is a wedge of circles and each Y_i is an immersed combinatorial circle, then $\langle X \mid \{Y_i\}_{i\in\mathcal{I}}\rangle$ is a group presentation in the usual sense (each $C(Y_i)$ is a disc) and $\mathrm{Cay}(X^*)$ is the associated Cayley graph of G. When X is a graph and each Y_i is an immersed graph, then the above cubical presentation is a graphical presentation in the sense of [RS87, Gro03, Oll06].

Remark 2.2 (Elevations)**.** Note that the local isometries $Y_i \to X$ lift to local isometries $Y_i \to \mathrm{Cay}(X^*)$ (in fact, under the small-cancellation conditions we shall soon be assuming, the latter maps are embeddings [Wis, Section 4]). We use the term *elevation* to refer to a lift $\widetilde{Y}_i \to \widetilde{X}$ of the map $\widetilde{Y}_i \to Y_i \to X$, where $\widetilde{Y}_i \to Y_i$ is the universal covering map. Since $Y_i \to X$ is a local isometry, it is π_1–injective and $\widetilde{Y}_i \to \widetilde{X}$ is a combinatorial embedding with convex image.

We now review some background about cubical small-cancellation theory, details of which can be found in [Wis].

Definition 2.3 (Abstract cone-piece, abstract wall-piece, cone-piece, wall-piece, piece)**.** Given a CAT(0) cube complex $\widetilde{X}$, and convex subcomplexes U, V, let $\mathrm{Proj}(U \to V)$ be the subcomplex of V defined as follows. First, a closed 1–cube e of V is in $\mathrm{Proj}(U \to V)$ if e is dual to a hyperplane intersecting U. Then add any cube of V whose 1–skeleton appears.

Let $\langle X \mid \{Y_i\}_{i\in\mathcal{I}}\rangle$ be a cubical presentation. Let $A, B \in \{Y_i\}_{i\in\mathcal{I}}$. An *abstract cone-piece* of B in A is a component of $\mathrm{Proj}(\widetilde{B} \to \widetilde{A})$, where $\widetilde{B}, \widetilde{A}$ are components of the preimages of B, A respectively, under the covering map $\widetilde{X} \to X$, satisfying $\widetilde{B} \neq \widetilde{A}$. If H is a hyperplane of X not intersecting B, then, likewise, an *abstract wall-piece* of H in B is a component of $\mathrm{Proj}(\mathcal{N}(\widetilde{H}) \to \widetilde{B})$, where $\mathcal{N}(\widetilde{H})$ denotes the *carrier* of $\widetilde{H}$. A *cone-piece* is a path in an abstract cone-piece, and a *wall-piece* is a path in an abstract wall-piece, and a *piece* is a path which is either a cone-piece or a wall-piece.

Definition 2.4 ($C'(\alpha)$ condition)**.** The cubical presentation $\langle X \mid \{Y_i\}_{i\in\mathcal{I}}\rangle$ satisfies the cubical $C'(\alpha)$ small-cancellation condition if $\mathrm{diam}(\mathcal{P}) < \alpha\|Y_i\|$ for all abstract pieces $\mathcal{P}$ and all $i \in \mathcal{I}$, where $\|Y_i\|$ denotes the infimum of the lengths of essential closed paths in Y_i. In this case, we say that $\langle X \mid \{Y_i\}_{i\in\mathcal{I}}\rangle$ is a $C'(\alpha)$ *presentation* and $G = \pi_1 X^*$ is a $C'(\alpha)$ *group*.

Given a cubical presentation $\langle X \mid \{Y_i\}_{i\in\mathcal{I}}\rangle$ and a closed path $P \to X$ that is nullhomotopic in X^*, van Kampen's lemma provides a disc diagram $(D, \partial_p D) \to (X^*, X)$ whose boundary path $\partial_p D = P$. (We always use the operator ∂_p, applied to a disc diagram, to denote the boundary path.) The 2–cells of such a diagram are either squares (mapping to 2–cubes of $X \subseteq X^*$) or 2–simplices mapping to cones over the various Y_i. Since P avoids cone-points, the 2–simplices of D are partitioned into classes: for each vertex of D mapping to a cone-point in X^*, the incident 2–simplices are arranged cyclically around the vertex to form a subspace C of D which is equal to the cone on its boundary path (a path in D mapping to X). The subspace C is a *cone-cell*. In practice, we ignore the subdivision of C into 2–simplices and regard C as a 2–cell of D.

The *complexity* of D is the pair (c, s), where c is the number of cone-cells and s is the number of squares. Taking the complexity in lexicographic order, we always consider diagrams $(D, \partial_p D) \to (X^*, X)$ which are *minimal* in the sense that the complexity of D is lexicographically minimal among diagrams with boundary path $\partial_p D$. This implies that for each cone-cell C, the path $\partial_p C \to D \to Y \to X$ is essential.

Remark 2.5 (Dual curves and hexagon moves)**.** Let $D \to X$ be a square diagram. A *dual curve* in D is a path which is the concatenation of midcubes of squares of D that starts and ends on $\partial_p D$, where a *midcube* of a square $[-\frac{1}{2}, \frac{1}{2}]^2$ is obtained by restricting exactly one coordinate to 0 and a midcube of a 1–cube is its midpoint. If X is a nonpositively-curved cube complex, then each dual curve maps to a hyperplane. If K is a dual curve in D, then the union of all closed cubes intersecting K is its *carrier*.

More generally, if $D \to X^*$ is a disc diagram, then one can define dual curves as above, but any dual curve has its two ends either on $\partial_p D$ or on the boundary path of a cone-cell of X^*.

A *hexagon move* is a homotopy of the diagram $D \to X^*$ that fixes the boundary path and the cone-cells and their boundary paths, while modifying the square part of D. Specifically, if s_1, s_2, s_3 are squares in D arranged cyclically around a central vertex v, forming a hexagonal subdiagram E of D, then X must contain a 3–cube c with a corner at the image of v formed by the images of s_1, s_2, s_3. The (hexagonal) boundary path of E maps to a combinatorial path in c, and we can replace E by a diagram E' formed from the other 3 squares on the boundary of c; this yields a new diagram $D' \to X^*$, with the same boundary path as D, formed by replacing E by E'. This modification is a hexagon move. Hexagon moves are used to reduce area in various ways; detailed accounts can be found in e.g. [Wis, Wis12].

Definition 2.6 (External cone-cell, internal cone-cell, internal path)**.** The cone-cell C of the disc diagram D is *external* if $\partial_p C = QS$, where Q is a non-trivial subpath of $\partial_p D$ (i.e. containing at least one 1–cell) and S is an *internal* path in the sense that no 1–cell of S lies on $\partial_p D$. The cone-cell C is *internal* if $\partial_p C$ and $\partial_p D$ have no common non-trivial subpath.

Remark 2.7 (Rectification, angling, and curvature)**.** Given a disc diagram $(D, \partial_p D) \to (X^*, X)$, one can *rectify* D, to produce a *rectified diagram* $\bar{D}$, by removing some internal open 1–cells, so that D is subdivided into cone-cells, *rectangles* which are obtained from square ladders by deleting the internal open 1–cells, and complementary regions called *shards*. The specific procedure for doing so is discussed in [Wis, Section 3.f], but we will not require the details here. After rectifying D, each corner in each of the resulting 2–cells is assigned an angle according to one of several possible schemes. Although we need not get into the details of the angle-assignments, we will need to assume that they follow the *split-angling* defined in [Wis]. As explained in [Wis, Section 3.g], this yields a notion of *curvature* in D concentrated at 0–cells and 2–cells.

Definition 2.8 (Generalized corner, spur, shell)**.** A *(positively-curved) shell* C in the disc diagram D is an external cone-cell whose curvature is positive; the boundary path of a shell has the form QS, where the *outer path* Q is a subpath of the boundary path of D, and the *inner path* S has no open 1–cell on $\partial_p D$. A *spur* in D is a vertex v in $\partial_p D$ so that the incoming and

outgoing 1–cells of $\partial_p D$ map to the same 1–cell of X, i.e. v is the second vertex in a subpath of $\partial_p D$ of the form ee^{-1}, where $e \to X$ is a 1–cell. A *generalized corner* is a path ef in D, where each of e, f is an edge, so that the dual curves emanating from e, f cross inside a square s of D, as shown in Figure 1.

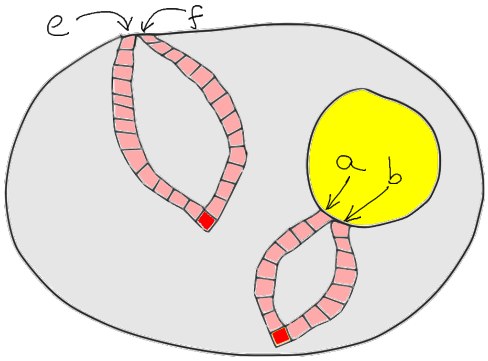

FIGURE 1. ef and ab are generalized corners of the shaded squares. ef lies on the boundary of D, while ab lies on the boundary of a cone-cell.

Remark 2.9 (Defects and curvature)**.** We now briefly review some notions related to curvature, from [Wis], that we will require below. Given a rectified disc diagram $\bar{D}$, we can assign an *angle* $\angle(c)$ – a real number – to each corner c of each 2–cell (i.e. to each 1–cell of each vertex-link). The *defect* $\mathfrak{d}(c)$ at the corner c is $\mathfrak{d}(c) = 2\pi - \angle(c)$. The *curvature* $\kappa(v)$ at a vertex v of $\bar{D}$ is $\kappa(v) = 2\pi - \sum \angle(c) - \pi\chi(\mathrm{Lk}(v))$, where $\mathrm{Lk}(v)$ is the link of v and the sum is taken over the 1–cells c of $\mathrm{Lk}(v)$. The *curvature* $\kappa(f)$ at a 2–cell f of $\bar{D}$ is $\kappa(f) = 2\pi - \sum \mathfrak{d}(c)$, where c varies over the corners of f.

We will need the following theorem, which follows immediately from the "combinatorial Gauss-Bonnet Theorem", i.e. Theorem 4.6 in [MW02]; this theorem dates, in one form or another, to [Bri48, Ger87, BB96].

Theorem 2.10 (Gauss-Bonnet for diagrams)**.** *Let $\bar{D} \to X^*$ be a rectified disc diagram. Then*

$$\sum_f \kappa(f) + \sum_v \kappa(v) = 2\pi,$$

where f varies over the 2–cells of $\bar{D}$ and v varies over the 0–cells of $\bar{D}$.

Remark 2.11 (Pushing generalized corners to the boundary)**.** If ef is a generalized corner of a square s, and ef lies along $\partial_p D$, and the subdiagram bounded by the carriers of the dual curves emanating from e, f is a square diagram, then we can perform a series of hexagon moves (see [Wis, Section 2]) to homotop D, fixing its boundary path, so that there is a square with boundary path $efe'f'$, i.e. we can move squares to the boundary. If, on the other hand, the same situation occurs except ef lies on the boundary of some cone-cell C mapping to a relator Y, then convexity of Y allows us to "absorb" the square s into C, lowering complexity of D.

Definition 2.12 (Padded ladder, ladder)**.** A *padded ladder* is a disc diagram $D \to X^*$ (or $\widetilde{X}^*$) with the following structure. First, there is a sequence $C_1, \dots, C_n$, where each C_i is a cone-cell or vertex of D, so that C_i, C_k lie in distinct components of $D - C_j$ whenever $i < j < k$. The diagram D is an alternating union of these vertices and cone-cells with a sequence of subdiagrams $R_0, \dots, R_n$ called *pseudorectangles*, so that:

(1) $\partial_p D$ is a concatenation $P_1 P_2^{-1}$, where each of P_1, P_2 starts on R_0 and ends on R_n.

(2) $P_1 = \nu_0\rho_0\alpha_1\rho_1\cdots\alpha_n\rho_n$ and $P_2 = \varrho_0\gamma_1\varrho_1\cdots\gamma_n\varrho_n\mu_{n+1}$.
(3) $\partial_p C_i = \mu_i\alpha_i\nu_i^{-1}\gamma_i^{-1}$.
(4) $\partial_p R_i = \nu_i\rho_i\mu_{i+1}^{-1}\varrho_i^{-1}$.
(5) Each R_i is a *square diagram*, i.e. contains no cone-cells.
(6) For each i, any dual curve in R_i emanating from ρ_i ends on ϱ_i and vice versa. Hence any dual curve emanating from ν_i ends on μ_{i+1} and vice versa.
(7) For each i, no two dual curves emanating from μ_i cross.

See Figure 2 for a picture illustrating the notation. We say that R_i is *horizontally degenerate* if $|\mu_{i+1}| = |\nu_i| = 0$ and *vertically degenerate* if $|\rho_i| = |\varrho_i| = 0$. If R_0, R_n are vertically degenerate,

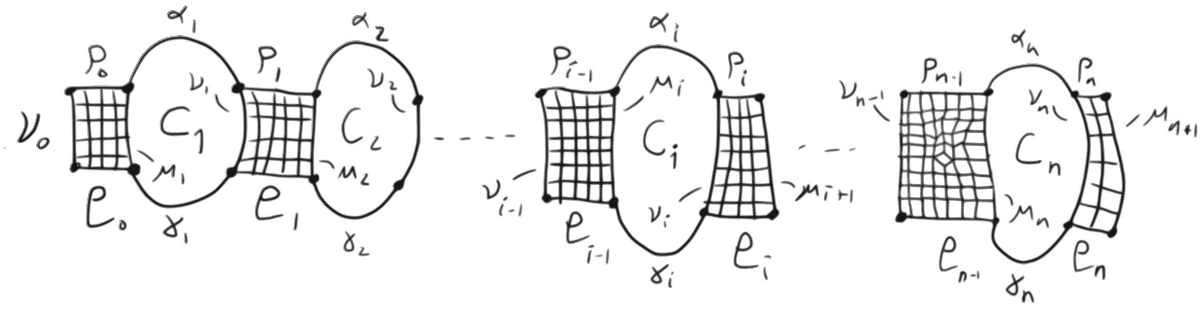

FIGURE 2. A padded ladder.

then D is a *ladder*. (A padded ladder is thus a special case of what Jankiewicz calls a *generalized ladder* in [Jan14], while the definition of ladder given here is equivalent to that in [Wis]).

We require the following three crucial facts, due to Wise [Wis]. These are tailored to our specific situation; the statements in [Wis] are more general. (That the cubical $C'(\frac{1}{12})$ condition suffices to make a cubical presentation "small-cancellation" in the sense required in the following theorems can be seen by tracing through various computations in [Wis, Section 3] or consulting the list of examples in Section 3.p of [Wis].)

Theorem 2.13 (Ladder theorem)**.** *Let $\langle X \mid \{Y_i\}\rangle$ be a cubical $C'(\frac{1}{12})$ presentation. Let $D \to X^*$ be a minimal disc diagram such that the corresponding rectified diagram has exactly two positively-curved cells along $\partial_p D$. Then D is a ladder.*

Theorem 2.14 (Greendlinger's lemma)**.** *Under the hypotheses of Theorem 2.13, if $D \to X^*$ is a minimal disc diagram, then either D consists of a single vertex or cone-cell, or D is a ladder, or D contains at least three shells or spurs or generalized corners along $\partial_p D$.*

The following, with $C'(\frac{1}{8})$ replaced by $C'(\frac{1}{12})$, is Lemma 3.46 of [Wis]; the improved constant follows from a result of Jankiewicz [Jan14].

Theorem 2.15 (Short inner paths)**.** *Let $\langle X \mid \{Y_i\}\rangle$ be a cubical $C'(\frac{1}{8})$ presentation. Let $D \to X^*$ be a disc diagram and let C be a shell in D with boundary path QS, with Q a maximal common subpath of $\partial_p C$ and $\partial_p D$, and S an internal path. Suppose that QS is essential in the relator Y to which C maps, and that S is of minimal length among all paths $S' \to Y$ that are homotopic rel endpoints in Y to S. Finally, suppose that the total curvature contribution from C is $< \pi$. Then $|S| < |Q|$.*

2.1. **The classification theorem.** Using the above fundamental results, we produce a cubical small-cancellation version of Strebel's classification of triangles in classical small-cancellation groups [Str90]. An *exposed square* in a disc diagram D is a square with two consecutive edges

on $\partial_p D$. A *tripod* is a triangle diagram with no cone-cells or squares. We can now state our classification of triangles:

Theorem 2.16 (Classification of triangles in cubical $C'(\frac{1}{144})$ groups)**.** *Let $\langle X \mid \{Y_i\}_{i\in\mathcal{I}}\rangle$ be a cubical presentation satisfying the $C'(\frac{1}{144})$ condition. Let $\alpha, \beta, \gamma \to \mathrm{Cay}(X^*)$ be combinatorial geodesics so that $\alpha\beta\gamma$ is a geodesic triangle. Then there exists a disc diagram $(D, \partial_p D) \to (X^*, X)$ with boundary path $\alpha'\beta'\gamma' \to \mathrm{Cay}(X^*) \to X^*$ lying in X, so that the following hold. First, $\alpha \to X$ and $\alpha' \to X$ co-bound a bigon $B \to X$ (i.e. they are square-homotopic) and the same is true of β, β' and γ, γ'. Second, D is of one of the following types.*

1. *(3–**shell generic**:) D has exactly three external cone-cells, C_1, C_2, C_3, respectively containing the points $\alpha \cap \beta, \beta \cap \gamma, \gamma \cap \alpha$. There is exactly one cone-cell M that intersects α, β, γ. Moreover, D is the union of three ladders, L_1, L_2, L_3 so that $L_i \cap L_j = M$ for all i, j. In particular, every cone-cell except M intersects exactly two of the geodesics α, β, γ.*
2. *(3–**shell tripod**:) D has exactly three external cone-cells, C_1, C_2, C_3, respectively containing the points $\alpha \cap \beta, \beta \cap \gamma, \gamma \cap \alpha$. Every other cone-cell intersects exactly two of the geodesics α, β, γ. In this case, D is the union of 3 (possibly padded) ladders L_1, L_2, L_3 and a tripod triangle $P_1P_2P_3 \to X$ so that L_i intersects the other two ladders in the path P_i.*
3. *(2–**shell generic**:) Same as 3–shell generic, except exactly one of C_1, C_2, C_3 is a spur or exposed square instead of a cone-cell.*
4. *(2–**shell tripod**:) Same as 3-shell tripod, except exactly one of C_1, C_2, C_3 is a spur or exposed square instead of a cone-cell.*
5. *(1–**shell generic**:) Same as 2–shell generic, except exactly two of C_1, C_2, C_3 are spurs or exposed squares.*
6. *(1–**shell tripod**:) Same as 2–shell tripod, except exactly two of C_1, C_2, C_3 are spurs or exposed squares.*
7. *(**No-shell generic**:) Same as 3–shell generic, except C_1, C_2, C_3 are all spurs or exposed squares.*
8. *(**No-shell tripod**:) Same as 3–shell tripod except C_1, C_2, C_3 are all spurs or exposed squares. This includes the case where $\alpha\beta\gamma$ is nullhomotopic in X, in which case D is a tripod.*
9. *(**Degenerate triangle**:) D is a single vertex or cone-cell, or D is a ladder. In this case, at least one of α, β, γ is trivial.*

The diagram $D \to X^*$ is a standard diagram *for the triangle $\alpha\beta\gamma$. The eight non-degenerate cases are shown in Figure 3.*

Remark 2.17 (Media and small-cancellation parameters)**.** Note that the standard diagram depends only on the endpoints of the geodesics α, β, γ. Just as it is usual in CAT(0) cube complexes to homotop geodesics, fixing their endpoints, in order to minimize the area of diagrams, here we are not married to particular geodesics, just to square-homotopy classes rel endpoints. In particular, if $\alpha\beta\gamma$ bounds a disc diagram in X, then D is a tripod. When $\mathcal{I} = \emptyset$, Theorem 2.16 just says: any three 0–cubes in a CAT(0) cube complex determine a geodesic tripod.

More generally, as illustrated by Figure 3, Theorem 2.16 should be interpreted as saying that the vertices of the triangle have a "median" which is either a vertex or a cone-cell, and there is a geodesic triangle connecting the given three points, each of whose sides passes within a wall-piece of the "median". In other words, given 0-cells $a, b, c \in \mathrm{Cay}(X^*)$, the "convex hulls" of the three possible pairs mutually coarsely intersect.

At the other extreme, when X is a graph and each Y_i is an immersed circle, Theorem 2.16 generalizes a weak version of Strebel's classification of triangles [Str90, Theorem 43]; specifically,

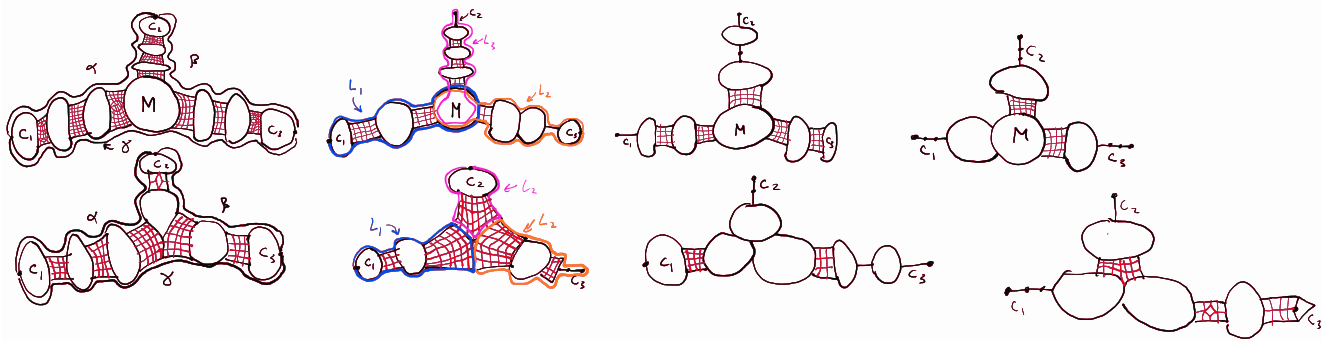

FIGURE 3. Clockwise from top left: cases (1),(3),(5),(7),(8),(6),(4),(2) from Theorem 2.16. The features of positive curvature are labelled, and the constituent ladders in some cases are labelled. Any of the spurs may instead be exposed squares and vice versa. The ladder case is not shown.

Theorem 2.16 provides the same classification as Strebel's result, but, because the proof must work in the more general context of cubical presentations, we require stronger metric small-cancellation conditions than Strebel needs in the classical setting.

Our need for $\frac{1}{144}$ rather than $\frac{1}{6}$ comes from our invocation of: the ladder theorem and Greendlinger lemma from [Wis], which require $C'(\frac{1}{12})$ condition (although Jankiewicz has improved this to $C'(\frac{1}{8})$ condition in the case of the Greendlinger lemma [Jan14]); the short inner paths condition, which requires $C'(\frac{1}{8})$; and Lemma 3.46 of [Wis], which requires $C'(\frac{1}{12})$. It appears that in the classical case, when there are no squares, each of these statements holds under the $C'(\frac{1}{6})$ condition. However, we use $C'(\frac{1}{144})$ to do appropriate curvature computations analogous to those in the proof of Theorem 3.20 of [Wis], to rule out internal cone-cells in diagrams. It seems as though one must make small modifications in order to have this part of the curvature-based proof work under $C'(\frac{1}{6})$ condition whenever there are no squares.

Proof of Theorem 2.16. This is essentially a meticulous application of Theorem 2.14, Theorem 2.13, and Theorem 2.15, following and followed by appropriately chosen square homotopies; the main point is a curvature computation to eliminate the possibility of internal cone-cells (in the Strebel classification in the case of classical $C'(\frac{1}{6})$ condition, one of the primary features is that the disc diagrams do not have internal cells). This computation, using a slightly modified version of the proof of Theorem 3.20 of [Wis], is why we need the $C'(\frac{1}{144})$ condition.

Choosing α', β', γ' and constructing D: Given a geodesic $P \to \mathrm{Cay}(X^*)$, let $[P]$ be the set of geodesics Q that have the same endpoints as P and the additional property that PQ^{-1} bounds a disc diagram containing no cone-cells, i.e. there is a disc diagram $E \to X$ whose boundary path is $PQ^{-1} \to \mathrm{Cay}(X^*) \to X^*$.

Choose a disc diagram $D \to \widetilde{X}^*$ so that $\partial_p D = \alpha'\beta'\gamma'$, where $\alpha' \in [\alpha], \beta' \in [\beta], \gamma' \in [\gamma]$. Choose D so that the complexity is minimal among all disc diagrams with the given properties.

Abusing notation slightly, we now temporarily regard D as the *rectified diagram* constructed in [Wis, Section 3]. This means that certain square ladders are regarded as single (rectangular) 2–cells, cone-cells are regarded as single cells, and the remaining parts of the diagram are 2–cells (formed by ignoring non-boundary 1–cells in certain square subdiagrams) called *shards*. Angles are assigned to corners according to the *split-angling* of [Wis, Section 3].

Applying the ladder theorem: By Theorem 2.13, either D is a ladder, so assertion (9) holds, and we are done, or D has at least 3 "features of positive curvature" – spurs, generalized corners, or shells – along $\partial_p D$. We assume the latter.

Applying the Greendlinger lemma: By Theorem 2.14, D has at least 3 features of positive curvature along the boundary, *each of which* is a shell, a spur, or a generalized corner.

We may assume that all generalized corners along $\partial_p D$ are actually squares with corners on $\partial_p D$. Indeed, let s be a square in D with a generalized corner on $\partial_p D$, i.e. the two dual curves K_1, K_2 intersecting s end at consecutive 1–cubes e_1, e_2 on $\partial_p D$. By a sequence of hexagon moves, we can modify D – without changing its boundary path or complexity – so that s lies along the boundary, i.e. e_1, e_2 are consecutive 1–cubes of s.

No positive curvature along geodesics (after square homotopy): Let s be a square of D so that $\partial_p s$ and $\partial_p D$ have a common subpath e_1e_2. For $i \in \{1, 2\}$, let e_i' be the 1–cube of s opposite e_i. If e_1e_2 is a subpath of one of the three constituent geodesics of $\partial_p D$ (say, α'), then we can modify α' in its square-homotopy class by replacing e_1e_2 by $e_2'e_1'$, resulting in a new diagram with the same number of cones and fewer squares. This contradicts our minimality assumption. Hence any square s with a corner on the boundary lies at the transition from α to β, or β to γ, or γ to α.

Now suppose that C is a positively-curved shell in D whose outer path P is a subpath of one of the named geodesics, say α and whose inner path we denote S. Suppose P is a subpath of one of the named geodesics bounding D, say α. Now, $\kappa(C) = 2\pi - \sum_c \mathfrak{d}(c)$, where c varies over the corners of C. The definition of the split-angling (see [Wis, Section 3]) says that $\mathfrak{d}(c) = 0$ at each corner c formed by a pair of 1–cells on $\partial_p C$ both lying on $\partial_p D$. Since α is a geodesic, it cannot be the case that $|S| = 0$, since otherwise we could replace P by S to shorten α. Thus, C has two corners, c_1, c_2, at the 0–cells where P, S meet. Each of these has angle $\pi/2$ by definition of the split-angling (see Figure 64 of [Wis]) and thus $\mathfrak{d}(c_1) = \mathfrak{d}(c_2) = \pi/2$. Thus, $\kappa(C) = \pi - \sum_c \mathfrak{d}(c)$, where c varies over the corners at non-endpoint vertices of S. Hence, if $\sum_c \mathfrak{d}(c) > 0$, the short inner paths condition, Theorem 2.15, shows that replacing P by S yields a strictly shorter path joining the endpoints of α, a contradiction. On the other hand, if $\sum_c \mathfrak{d}(c) = 0$, then Lemma 3.46 of [Wis] implies that S can be written as the concatenation of at most 7 pieces, and the small-cancellation assumption then implies that $|S| < |P|$, contradicting that α is a geodesic. Hence any such shell C has outer path $P = AB$, where A is a nontrivial terminal subpath of α, β, or γ and B is a nontrivial initial subpath of β, γ, or α.

Finally, since α, β, γ are geodesic, a spur of the form ee^{-1} cannot occur along each of α, β, γ so the only spurs consist of overlaps between α, β or β, γ or γ, α.

Hence, D has exactly three features of positive curvature along the boundary, which are subdiagrams C_1, C_2, C_3. For each i, $\partial_p C_i = OI$, where O is a subpath of $\partial_p D$ and I is an internal path, and O has at least one 1-cube on each of two distinct subpaths α, β, γ of the boundary.

No internal cone-cells: Let C be a cone-cell of D. Recall that C is *internal* if its boundary path intersects $\partial_p D$ in a set containing no 1–cube. By slightly modifying the "quick $\frac{1}{24}$" part of the proof of Theorem 3.20 of [Wis], we see that our small-cancellation condition is strictly stronger than the $C'(\frac{1}{72})$ condition required to ensure that each internal cone-cell contributes $< -4\pi$ of curvature, and all remaining internal features of the diagram contribute nonpositive curvature. Suppose that there are $n \geqslant 0$ internal cone-cells.

Let v be a 0–cube of D. Then the curvature contribution from v is:

(1) $\leqslant 0$ if v is internal or not contained in a 2–cell and not a spur;
(2) π if v is a spur;
(3) $\frac{\pi}{2}$ if v is the corner of a square along $\partial_p D$.

(Because we are in the rectified diagram, natural candidates for internal features of positive curvature, i.e. 0–cubes with three incident cyclically-arranged squares – have been relegated to the insides of shards, and do not actually contribute any positive curvature in the split-angling.)

Let f be a 2–cell of D (a cone-cell, rectangle, or a *shard* of the corresponding *rectified diagram* [Wis]). Then the curvature contribution is:

(1) $\leqslant 0$ if f is a rectangle or shard;
(2) $< -4\pi$ if f is an internal cone-cell;
(3) at most 2π if f is a shell.

Hence, our three features of positive curvature contribute a total of at most 6π of curvature, while the sum of the remaining curvatures is $< -4n\pi$. This contradicts Theorem 2.10 unless $n = 0$. Hence there are no internal cone-cells.

No shortly-external cone-cells: A cone-cell C in D is *shortly external* if its boundary path has the form QI, where I is internal and Q is a subpath of α, β, or γ. Note that $|Q| \leqslant |I|$ since α, β, γ are geodesics. Hence, by $C'(\frac{1}{144})$, the path I contains more than $72 = 3 \cdot 24$ transitions between pieces and thus the total angle-defect along I is more than 6π, so that the curvature contribution from C is less than -4π, so, as above, the Gauss-Bonnet theorem ensures that there are no shortly external cells in D.

At this point, we have completed the curvature computations in the proof, and now regard D as an ordinary (not rectified) diagram. This amounts to filling in the shards and rectangles with their constituent squares as in the original diagram.

Analysis of the cone-cells: Let C be a cone-cell in D. Suppose that for some $\delta \in \{\alpha, \beta, \gamma\}$, there is a subpath $\delta' = PQR$ of δ, where P, R are subpaths of the boundary path of C, and the terminal vertex of P and initial vertex of R subtend a subpath Q' of $\partial_p C$, so that QQ' is a subdiagram E of D between C and $\partial_p D$. If E contains no cone-cell, then E is a square diagram between the relator Y_i to which C maps and the geodesic Q, so by local convexity of Y_i in X, we have that $E \to X$ factors through $Y_i \to X$. Hence E could have been absorbed into the cone-cell C, whence minimality of the complexity of D ensures that E is trivial, i.e. $Q = (Q')^{-1}$. Moreover, we may assume that Q, Q' have no common 1–cell, by considering a minimal example.

Thus, assume C is innermost, so that any cone-cell C_0 in E embeds and has connected intersection with Q, and assume that E contains such a cone-cell C_0. Then C_0 is not internal in E, for then it would be internal in D. Moreover, C_0 cannot be a shell with outer path on Q, because then it would either be an already considered illegal feature of positive curvature or a shortly-external cone-cell in D. If C_0 is a shell, it must therefore have outer path (within E) of the form TU, where T is a terminal subpath of Q and U an initial subpath of Q'. But then C_0 is shortly-external in D, which is impossible. Hence any cone-cell C_0 in E has boundary path of the form $TI_1U_1I_2U_2 \cdots I_kU_kI_{k+1}$, where each I_i is internal to E, each U_i is a subpath of Q', and T is a subpath of Q. But then C_0 is either a positively-curved shell in D or shortly-external in D, neither of which is possible. Thus E is a square diagram, which was dealt with above.

We conclude that for each cone-cell C of D, the path $\partial_p C$ has connected intersection with each of α, β, and γ. Moreover, $\partial_p C$ intersects at least two of the paths α, β, γ.

Dividing into cases: We are now in the following situation: D is not a ladder, and has precisely 3 features of positive curvature along its boundary path, which are subdiagrams C_1, C_2, C_3. The subdiagram C_1 has boundary path ABI, where A is a nontrivial terminal subpath of α, B is a nontrivial initial subpath of β, and I is a (possibly trivial) path. Moreover, either C_1 is a single cone-cell (a shell) or C_1 is a spur, $|I| = 0$, and A is an edge and $B = A^{-1}$, or C_1 is a square, $|A|, |B| \geqslant 1$, and $|I| \leqslant 2$. The same description holds for C_2 (with β, γ replacing α, β) and C_3 (with γ, α replacing α, β).

Moreover, every cone-cell C of D not in $\{C_1, C_2, C_3\}$ has connected intersection with each of α, β, γ and intersects at least 2 of these paths. We call C a *median-cell* if C actually intersects all three of these paths, and a *tail-cell* otherwise. We emphasise that if C is a median cell, then it has nonempty, *connected* intersection with each of α, β, γ. Hence, if C is a median-cell,

then C separates D into three complementary components, each disjoint from one of the paths α, β, γ. Hence C is the unique cone-cell of D intersecting each of α, β, γ. (We note that there may be other disc diagrams with the same boundary path, containing a different median-cell.)

We now divide into cases. First, if D contains a median-cell, then we are in one of the *generic* cases, i.e. we will show that one of (7),(5),(3),(1) holds, according to how many of $\{C_1, C_2, C_3\}$ are spurs or shells. Otherwise, we will show that one of (8),(6),(4),(2) holds. This will complete the proof.

The generic cases: Suppose that D has a (unique) median-cell M and let $i \in \{1, 2, 3\}$. Let $\delta, \delta' \in \{\alpha, \beta, \gamma\}$ be the parts of the boundary path of D that intersect C_i. Let $\partial_p M = AP_1BP_2CP_3$, where A, B, C are respectively subpaths of α, β, γ and P_1, P_2, P_3 are internal paths. Write $\alpha = \bar{\alpha}A\check{\alpha}, \beta = \bar{\beta}B\check{\beta}, \gamma = \bar{\gamma}C\check{\gamma}$. Consider the subdiagram L_1 bounded by $A\check{\alpha}\bar{\beta}BP_2CP_3$. The ladder theorem (Theorem 2.13) and our above analysis of the possible features of positive curvature in (the rectification of) D shows that L_1 is a ladder. The ladders L_2, L_3 are constructed analogously.

The tripod cases: Suppose there is no median. Then we have a subdiagram T of D with boundary path $AP_1BP_2CP_3$, where A is a subpath of α, B a subpath of β, C a subpath of γ, and P_1, P_2, P_3 internal subpaths that lie on innermost cone-cells in D or, if they do not exist, spurs or exposed squares in $\{C_1, C_2, C_3\}$. By construction, T is a possibly degenerate square diagram, and by convexity of relators and minimality, for each path $Q \in \{A, B, C, P_1, P_2, P_3\}$, no two dual curves in T emanating from Q can cross. Moreover, no dual curve travels from Q to Q or to the next named subpath, for otherwise we could reduce complexity. Some possibilities are shown in Figure 4.

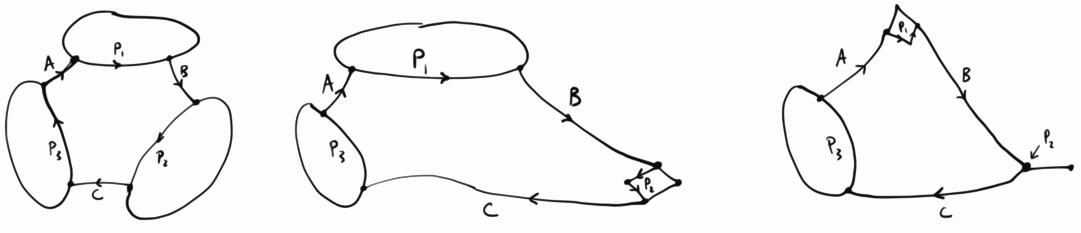

FIGURE 4. Some possibilities for the internal square subdiagram in the tripod cases.

For convenience, we lift T to a diagram $T \to \widetilde{X}$ (the CAT(0) cube complex $\widetilde{X}$, not the generalized Cayley graph). Here, analysis of the dual curves shows that T decomposes as required; the analysis is indicated in Figure 5. First, consider dual curves in T traveling from A

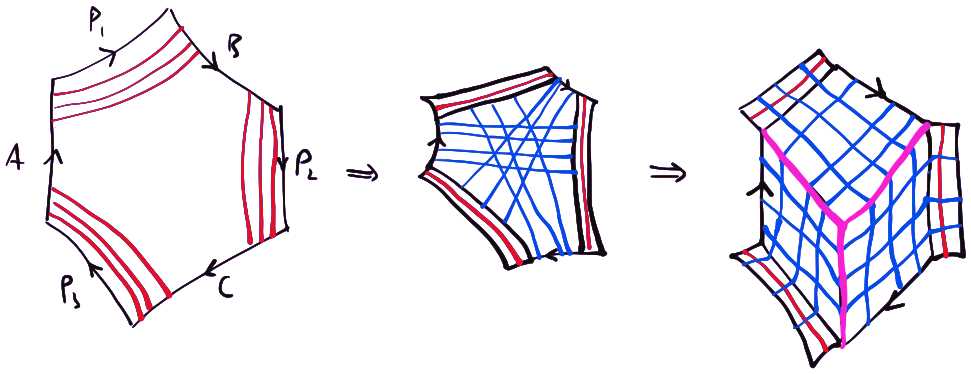

FIGURE 5. The final square diagram analysis in the tripod cases.

to B, B to C, or C to A. Taking the union of all carriers of such dual curves yields rectangles attached to P_1, P_2, P_3. Now consider the subdiagram that remains. It is a hexagon bounded by subpaths of A, B, C and parts of carriers of dual curves. Dual curves in the subdiagram must travel from a subpath of A, B, or C to the antipodal dual-curve carrier. Dual curves emanating from the same "syllable" of the boundary path do not cross, and we conclude, as at right in Figure 5, that this subdiagram is a "corner of a subdivided cube". It is now easy to deduce the padded ladder decomposition of D. (Various parts of the picture may be degenerate, as suggested in Figure 4.) □

3. Detecting WPD elements using the classification of triangles

In this section, we make the following assumptions and conventions:

(1) $\langle X \mid \{Y_i\}_{i\in\mathcal{I}}\rangle$ is a cubical presentation satisfying the $C'(\frac{1}{144})$ condition, X^* is the presentation complex, and $\widetilde{X}^*$ is the universal cover. We assume that X is locally finite, although we do **not** assume X is **uniformly** locally finite. For example, X can contain finite cubes of arbitrarily large dimension.
(2) Denote by d the graph metric on $\mathrm{Cay}(X^*)^{(1)}$, and by $\mathsf{d}_{\widetilde{X}}$ the graph metric on $\widetilde{X}^{(1)}$.
(3) Let there be a δ–hyperbolic space $\mathcal{H}$ and a map $\Pi : \mathrm{Cay}(X^*) \to 2^{\mathcal{H}}$ so that:
   (a) There exists $\mathfrak{e} \geqslant 0$ so that $\mathrm{diam}(\Pi(x)) \leqslant \mathfrak{e}$ for all $x \in \mathrm{Cay}(X^*)$.
   (b) There exists $\mathfrak{f} \geqslant 0$ so that $\mathsf{d}_{\mathcal{H}}(\Pi(x), \Pi(y)) \leqslant \mathfrak{f}\mathsf{d}(x, y) + \mathfrak{f}$ whenever $x, y \in \mathrm{Cay}(X^*)^{(0)}$ are not cone-points.
   (c) If $Y_i \subset \mathrm{Cay}(X^*)$ is any relator, then $\mathrm{diam}(\Pi(Y_i)) \leqslant \mathfrak{e}$.
   (d) $\pi_1 X^*$ acts by isometries on $\mathcal{H}$ in such a way that Π is $\pi_1 X^*$–equivariant.
   (e) Let $\{I_j\}$ be a set of combinatorial intervals, each of positive length, so that there is a cubical immersion $\prod_j I_j \to X$, lifting to a cubical isometric embedding $\iota : \prod_j I_j \to \mathrm{Cay}(X^*)$. Then $\mathrm{diam}(\mathrm{im}(\Pi \circ \iota)) \leqslant \mathfrak{e}$.

Under these conditions, we will prove a lemma showing that $\pi_1 X^*$ contains a WPD isometry of $\mathcal{H}$. Later, we choose specific $\mathcal{H}$ and Π allowing the lemma to be invoked.

Lemma 3.1 (Ladders are thin between cone-cells)**.** *Let $L \to \widetilde{X}$ be a padded ladder with boundary path $\alpha\beta\gamma$, where $\alpha, \beta : [0, n] \to \mathrm{Cay}(X^*)$ are geodesics and γ is a piece. Then there exists a uniform constant κ_0 so that for all $t \leqslant n$, either $\mathsf{d}(\alpha(t), \beta(t)) \leqslant \kappa_0$ or $\alpha(t), \beta(t)$ lie on the boundary path of a common cone-cell in L.*

Proof. Write $\alpha = \alpha_0\eta_1\alpha_1 \cdots \eta_n\alpha_n$ and $\beta = \beta_0\eta'_1 \cdots \eta'_n\beta_n$, where each α_i, β_i lies on the top or bottom boundary path of one of the constituent pseudorectangles of L and each η_i, η'_i lies on the boundary path of a cone-cell or cut-vertex. Denote by p_i the maximal piece in L between the i^{th} cone-cell or cut-vertex and the i^{th} pseudorectangle, so that the i^{th} pseudorectangle is bounded by $\alpha_i p_{i+1}^{-1} \beta_i^{-1} p_i$. See Figure 6, at left.

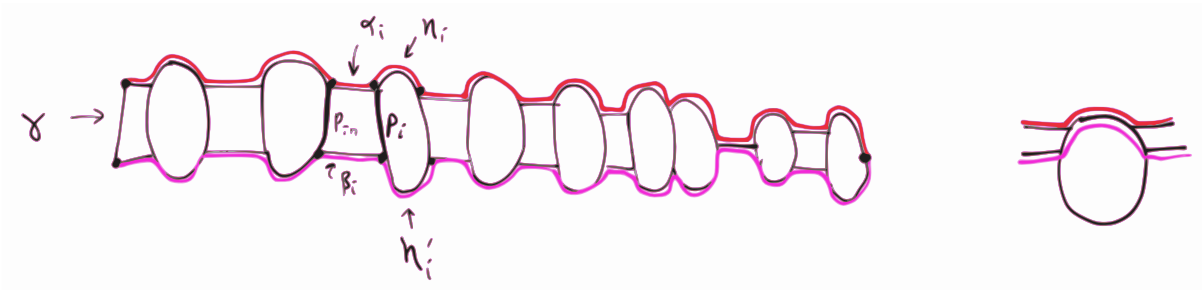

Figure 6. Ladders are thin relative to cone-cells.

By the small-cancellation conditions, there is a uniform M so that $|p_i| \leqslant M$ for all i. Since α, β are geodesics, we have $||\eta_i| - |\eta_i'|| \leqslant 2M$ for all i, for otherwise we could construct shortcuts, as shown at right in Figure 6. The lemma now follows easily. □

As usual (see e.g. [CS11]), $\tilde{g} \in \pi_1 X$ is *rank-one* if it is hyperbolic on $\widetilde{X}$ and none of its axes lies in an isometrically embedded Euclidean half-plane.

Lemma 3.2. *Let $\tilde{g} \in \pi_1 X$ act hyperbolically on $\widetilde{X}$, and suppose that $\tilde{g}$ is rank-one. Then for each $\tilde{x} \in \widetilde{X}^{(0)}$, there exists κ_1 so that the following holds: if $n \geqslant 0$ and $P, Q : [0, d] \to \widetilde{X}$ are combinatorial geodesics joining $\tilde{x}, g^n\tilde{x}$, then $\mathsf{d}_{\widetilde{X}}(P(t), Q(t)) \leqslant \kappa_1$ for $0 \leqslant t \leqslant d$.*

Proof. Let $\alpha \to \widetilde{X}$ be a combinatorial geodesic axis for $\tilde{g}$ and let $\tilde{a} \in \alpha$ be a 0–cube. Given $n \geqslant 0$, let $P, Q : [0, d_n] \to \widetilde{X}^{(1)}$ (where $d_n = \mathsf{d}_{\widetilde{X}}(\tilde{x}, g^n\tilde{x})$) be combinatorial geodesics joining $\tilde{x}, g^n\tilde{x}$ and let $D \to \widetilde{X}$ be a minimal-area disc diagram with boundary path PQ^{-1}. Note that $D \to \widetilde{X}$ is actually an isometric embedding on the 1–skeleton. Indeed, every dual curve in D travels from P to Q since P, Q are geodesics. Hence each dual curve maps to a distinct hyperplane, so that for any vertices $v, v' \in D$, the number of dual curves of D separating v, v' is equal to the number of hyperplanes in $\widetilde{X}$ separating their images.

Fix $t \in \{0, 1, \dots, d_n\}$. The above discussion shows that $\mathsf{d}_{\widetilde{X}}(P(t), Q(t))$ is bounded by the number of dual curves in D that travel from $P([0, t))$ to $Q((t, d_n])$, plus the number of dual curves from $Q([0, t))$ to $P((t, d_n])$. Each dual curve of the former type crosses each dual curve of the latter type. If no κ_1 with the claimed property exists, then since all but at most $2\mathsf{d}_{\widetilde{X}}(\tilde{a}, \tilde{x})$ hyperplanes that cross any given P_n cross α, we have that for all N, there are two sets $\mathfrak{H}, \mathfrak{V}$ of hyperplanes, each of cardinality N, so that $H \cap \alpha \neq \emptyset$ for all $H \in \mathfrak{H} \cup \mathfrak{V}$ and such that $H \cap V \neq \emptyset$ whenever $H \in \mathfrak{H}, V \in \mathfrak{V}$. This contradicts that $\tilde{g}$ is rank-one, since it implies that the convex hull of α contains 0–cubes arbitrarily far from α. □

Lemma 3.3 (Loxodromic implies WPD)**.** *Suppose $g \in \pi_1 X^*$ acts loxodromically on $\mathcal{H}$. Then for all $\epsilon > 0, x \in \mathcal{H}$, there exists $R \in \mathbb{N}$ so that*

$$|\{h \in G \mid \mathsf{d}_{\mathcal{H}}(hx, x) \leqslant \epsilon, \mathsf{d}_{\mathcal{H}}(hg^R x, g^R x) \leqslant \epsilon\}| < \infty,$$

i.e. g is a WPD element.

Proof. Fix $\epsilon > 0$ and let $x \in \mathcal{H}$; since $\mathcal{H}$ is Hausdorff-close to $\mathrm{Cay}(X^*) \subset \mathcal{H}$, we can assume x is a vertex of $\mathrm{Cay}(X^*)$. Fix $R \geqslant 10(\epsilon + \delta + \mathfrak{e})/\tau$, where $\tau \geqslant 1$ is the translation length of g on $\mathcal{H}$, and let $y = g^R x$. Let $h \in \pi_1 X^*$ satisfy $\mathsf{d}_{\mathcal{H}}(\Pi(x), \Pi(hx)) < \epsilon, \mathsf{d}_{\mathcal{H}}(\Pi(y), \Pi(hy)) < \epsilon$.

Let α be a combinatorial $\mathrm{Cay}(X^*)$–geodesic from y to x, let β be a $\mathrm{Cay}(X^*)$–geodesic from x to hx, let η be a geodesic from hy to y, and let γ be a geodesic from hx to y, so that we have geodesic triangles $\alpha\beta\gamma$ and $\eta(h\alpha)^{-1}\gamma$ with common side γ.

Applying the classification of triangles: By Theorem 2.16, we have a minimal disc diagram $D = D_1 \cup_\gamma D_2 \to \widetilde{X}^*$, with boundary path $\alpha\beta(h\alpha)^{-1}\eta$, with the following structure:

- D_1 has boundary path $\alpha\beta\gamma$ and D_2 has boundary path $\eta(h\alpha)^{-1}\gamma$.
- For $i \in \{1, 2\}$, the diagram D_i decomposes as $B_i^1 \cup B_i^2 \cup B_i^3 \cup S_i$, where S_i is a standard diagram in the sense of Theorem 2.16 and each B_i^j is a bigon diagram in X (i.e., no cone cells). The boundary path of S_i is a geodesic triangle $A_iB_iC_i$, where $A_1\alpha^{-1}$, $B_1\beta^{-1}, C_1\gamma^{-1}$ are the boundary paths of B_1^1, B_1^2, B_1^3 respectively, and $A_2(h\alpha), B_2\eta^{-1}$, and $C_2\gamma^{-1}$ are the boundary paths of B_2^1, B_2^2, B_2^3 respectively.
- S_i contains a constituent padded ladder L_i whose image in $\widetilde{X}^*$ projects under Π to a set of diameter at least $R - 2(\epsilon + \delta + \mathfrak{e})$, along with two ladders projecting to sets of diameter $\leqslant 10(\epsilon + \delta + \mathfrak{e})$ (which will play no role in the remainder of the proof). Specifically, the padded ladder L_1 is the subdiagram of D_1 obtained as follows: either D_1 is a ladder, in which case $L_1 = D_1$, or there is a cone-cell or tripod with 3 complementary components,

all of whose closures are padded ladders; L_1 is the padded ladder among these that contains y. The padded ladder L_2 is defined analogously.

The above notation is summarized in Figure 7.

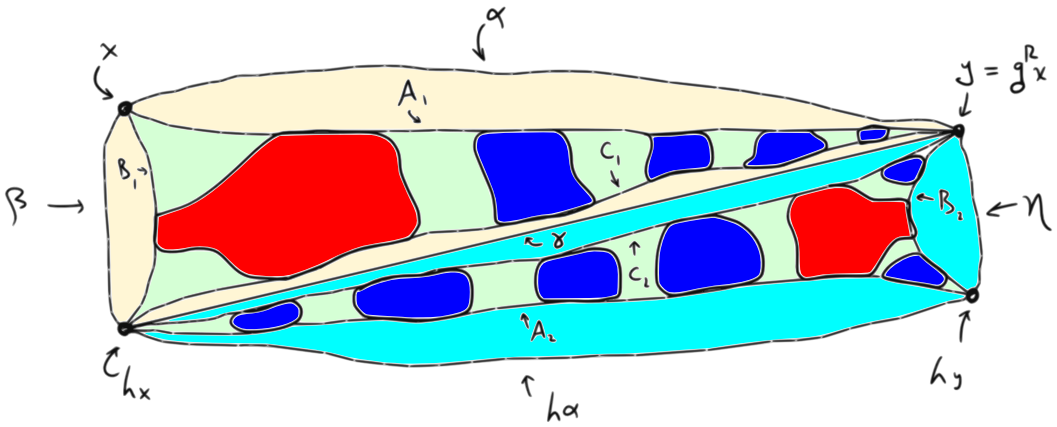

FIGURE 7. The two triangles in the proof of Lemma 3.3. The padded ladder L_1 is the subdiagram of S_1 between the red subdiagram (cone-cell or union of 3 square grids) intersecting A_1, B_1, C_1 and the point y. The ladder L_2 is the subdiagram of S_2 between the analogous red subdiagram (at right) and hx.

Consequences of axiality of g: Since g acts loxodromically on $\mathcal{H}$ and since $\Pi(A_1), \Pi(\alpha)$ fellow-travel in $\mathcal{H}$ (so $\Pi \circ A_1$ is a quasi-isometric embedding with constants independent of R, ϵ), the following hold, where A_1' is the part of A_1 on the boundary path of the ladder L_1:

- $A_1' = \rho_0\sigma_1\rho_1 \cdots \sigma_s\rho_s$, where each ρ_i lies on a pseudorectangle and each σ_i lies on the boundary path of a cone-cell;
- $|\sigma_i|, |\rho_i| \leqslant \Delta$, where Δ depends only on g and x, with the following exception: we may have $|\rho_i| > \Delta$ if the pseudorectangle carrying ρ_i is horizontally degenerate.

Similarly, the maximal subpath A_2' of A_2 lying on the ladder L_2 decomposes as $\varrho_0\varsigma_1 \cdots \varsigma_\ell\varrho_\ell$, where each ϱ_i lies on a pseudorectangle, each ς_i lies on a cone-cell, and each $|\varsigma_i|, |\varrho_i| \leqslant \Delta$, except that we may have $|\varrho_i| > \Delta$ if ϱ_i is carried on a horizontally degenerate pseudorectangle. See Figure 8.

For each i, let R_i be the pseudorectangle carrying ρ_i and let ρ_i' be the part of the boundary path of R_i parallel to (i.e. crossing the same dual curves as) ρ_i. Let K_i be the cone-cell carrying σ_i and let σ_i' be the part of $\partial_p K_i$ between ρ_{i-1}' and ρ_i', as shown in Figure 8. Let $C_1' = \rho_0'\sigma_1'\rho_1' \cdots \sigma_s'\rho_s'$ be the part of C_1 formed by concatenating these paths. Define ς_i', ϱ_i', and the resulting subpath C_2' of C_2 analogously.

Conclusion: Arguing as in the proof of Lemma 3.2 shows that C_1', C_2' synchronously fellow-travel at distance ξ_1 depending on Δ. Lemma 3.1, which shows that the pseudorectangles between ρ_i, ρ_i' and between ϱ_i, ϱ_i' are uniformly thin, shows that if $t_2 - t_1 > 2\Delta$, then there exists $t \in [t_1, t_2]$ so that $\mathsf{d}(A_1'(t), C_1'(t)) \leqslant \xi_2$, where ξ_2 is a uniform constant coming from the small-cancellation condition, and the same holds for A_2', C_2'. A reparameterization computation then shows that there exists $\xi_3 = \xi_3(\kappa_0, \kappa_1, \epsilon, \Delta)$, independent of x, y, so $\mathsf{d}(A_1(t), A_2(t)) \leqslant \xi_3$ for some t. Lemma 3.2 then produces a uniform ξ so that $\mathsf{d}(z, hz) \leqslant \xi$, where $z = \alpha(t)$. Since $\mathrm{Cay}(X^*)$ is locally finite and $\pi_1 X^*$ acts freely on $\mathrm{Cay}(X^*)$, the action of $\pi_1 X^*$ on $\mathrm{Cay}(X^*)$ is metrically proper and hence there are only finitely many such h, as claimed. □

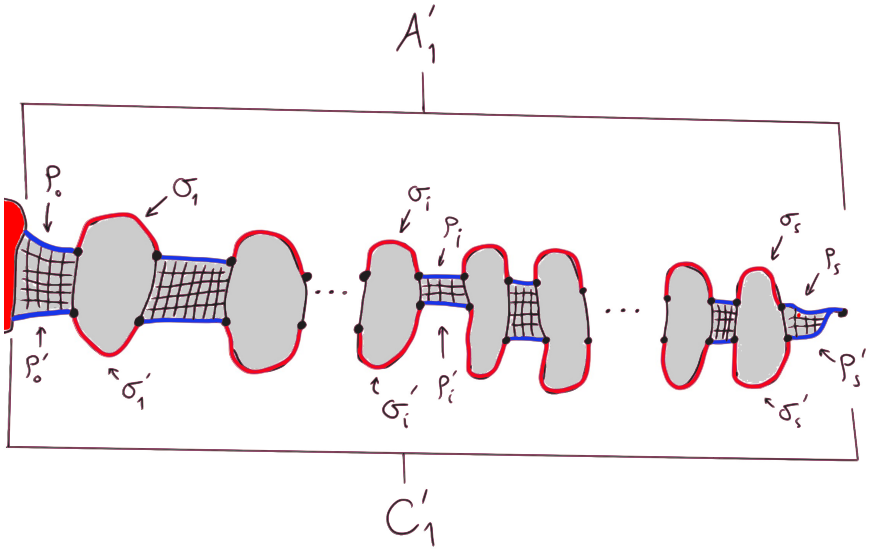

FIGURE 8. The ladder L_1.

4. The hyperbolic space $\mathcal{H}$ and the coarse projection

Let $\langle X \mid \{Y_i\}\rangle$ be a cubical $C'(\frac{1}{144})$ presentation and define a space $\mathcal{H}$ as follows. First, let $\mathcal{H}'$ be the 1–skeleton of $\widetilde{X}^*$. This consists of the 1–skeleton of $\mathrm{Cay}(X^*)$, together with a combinatorial cone on each lift of each Y_i. We form $\mathcal{H}$ from $\mathcal{H}'$ by adding a combinatorial cone on the carrier of each hyperplane.

We also have a projection $\Pi\colon \mathrm{Cay}(X^*) \to \mathcal{H}$, defined as follows. On the 1–skeleton of $\mathrm{Cay}(X^*)$, we declare Π to be the inclusion. If c is a cube of $\mathrm{Cay}(X^*)$ with $\dim(c) \geqslant 2$, we send c arbitrarily to a point in the image of its 1–skeleton. However, we require this choice to be made $\pi_1 X^*$–equivariantly, so that Π is $\pi_1 X^*$–equivariant. Obviously Π is 1–Lipschitz on the 1–skeleton of $\mathrm{Cay}(X^*)$. By construction, Π sends each cone to a set of diameter $\leqslant 2$, while cubical product subcomplexes of $\mathrm{Cay}(X^*)$ are sent to subsets of $\mathcal{H}$ with diameter at most 4. Hence, to see that $\mathcal{H}$ and Π satisfy the conditions required in Section 3, we need only to prove that $\mathcal{H}$ is hyperbolic.

Lemma 4.1 (Square bigons have thin projection)**.** *Let $\alpha, \beta \to \mathrm{Cay}(X^*)$ be geodesics with common endpoints, and suppose that $\alpha\beta$ bounds a disc diagram $D \to \widetilde{X}^*$ that does not contain any cone-cells. Then $\Pi(\alpha), \Pi(\beta)$ lie at uniformly bounded Hausdorff distance in $\mathcal{H}$.*

Proof. Let e be a 1–cube of α and let K be the dual curve in D emanating from e and mapping to a hyperplane H of $\mathrm{Cay}(X^*)$. Since α is a geodesic, K terminates at a 1–cube f of β, whence $\mathsf{d}_{\mathcal{H}}(\Pi(e), \Pi(f)) \leqslant 2$. Hence $\Pi(\alpha) \subseteq \mathcal{N}_2(\Pi(\beta))$, the closed 2-neighborhood of $\Pi(\beta)$, and the proof is complete by symmetry. □

Proposition 4.2. *There exists $\delta \geqslant 0$ so that the graph $\mathcal{H}$ is δ–hyperbolic.*

Proof. It suffices to prove that the 0–skeleton of $\mathrm{Cay}(X^*)$, with the subspace metric inherited from $\mathcal{H}$, is hyperbolic. First, suppose that $\alpha\beta\gamma$ is a geodesic triangle in $\mathrm{Cay}(X^*)$. Then Lemma 4.1 combines with Theorem 2.16 and the fact that Π sends cones to uniformly bounded sets to show that each of $\Pi(\alpha), \Pi(\beta), \Pi(\gamma)$ is contained in the δ'–neighborhood in $\mathcal{H}$ of the union of the other two, for some uniform δ'. The Guessing Geodesics Lemma (see e.g. [Ham07, Proposition 3.5][Bow14, Proposition 3.1]) now implies that $\mathcal{H}$ is δ–hyperbolic for some δ. □

Theorem 4.3. *Let $\langle X \mid \{Y_i\}_{i\in\mathcal{I}}\rangle$ be a $C'(\frac{1}{144})$ presentation with X locally finite, and let $G = \pi_1 X^*$. Then any $g \in G$ acting loxodromically on the space $\mathcal{H}$ constructed above acts on $\mathcal{H}$ as a WPD element, whence either G is virtually cyclic or acylindrically hyperbolic.*

Proof. The assertion that g is a WPD element follows from Lemma 3.3; hyperbolicity of $\mathcal{H}$ comes from Proposition 4.2. Applying [Osi15a, Theorem 1.2.$(AH_3 \Rightarrow AH_2)$] completes the proof. □

5. Proof of Theorem B

5.1. **Preservation of loxodromics.** In this section, we study the question of when G contains a loxodromic isometry of $\mathcal{H}$, using knowledge of which elements of $\pi_1 X$ act loxodromically on the *contact graph* $\mathcal{C}\widetilde{X}$ of $\widetilde{X}$, which is the intersection graph of the set of hyperplane-carriers in $\widetilde{X}$. Let $p : \widetilde{X} \to \mathrm{Cay}(X^*)$ be the obvious covering map (regarding $\mathrm{Cay}(X^*)$ as the cover of X corresponding to the subgroup $K = \langle\langle\{\pi_1 Y_i\}_{i\in\mathcal{I}}\rangle\rangle$ of $\pi_1 X$). Let $\widetilde{\mathcal{H}}$ be the graph obtained from $\widetilde{X}^{(1)}$ by coning off the 1–skeleton of each hyperplane-carrier.

Form a new graph $\widehat{\mathcal{H}}$ from $\widetilde{\mathcal{H}}$ by coning off every subgraph of $\widetilde{X}^{(1)} \subset \widetilde{\mathcal{H}}$ which is the 1–skeleton of an elevation $\widetilde{Y}_i \hookrightarrow \widetilde{X}$ of some $Y_i \to X$. Observe that p induces a quotient map $\hat{p} : \widehat{\mathcal{H}} \to \mathcal{H}$, which restricts to p on $\widetilde{X}^{(1)}$ and which sends the cone-point v_H over the hyperplane-carrier $\mathcal{N}(H)$ to the cone-point $v_{p(H)}$ over the hyperplane-carrier $p(\mathcal{N}(H))$. The map p also sends the cone over each $\widetilde{Y}_i$ to the cone over the corresponding lift $Y_i \to \widetilde{X}^*$ of $Y_i \to X$.

Remark 5.1. It is not hard to adapt the proof of Proposition 3.1 of [BHS14] to show that each $\widetilde{Y}_i^{(1)}$ is (uniformly) quasiconvex in the hyperbolic graph $\widetilde{\mathcal{H}}$. In conjunction with the small-cancellation assumption, this shows that $\widehat{\mathcal{H}}$ is again hyperbolic. We will not require this fact.

Lemma 5.2. *Let Γ be the graph with a vertex for each hyperplane-carrier in $\widetilde{X}$, and a vertex for each lift of each $\widetilde{Y}_i$ to $\widetilde{X}$, with adjacency corresponding to intersection. Then Γ is $\pi_1 X$–equivariantly quasi-isometric to $\widehat{\mathcal{H}}$ and $\widetilde{\mathcal{H}}$ is $\pi_1 X$–equivariantly quasi-isometric to $\mathcal{C}\widetilde{X}$.*

Proof. This is proved in [Hag14, Section 5] in the case where $\mathcal{I} = \emptyset$, so that Γ is the contact graph of $\widetilde{X}$ and $\widetilde{\mathcal{H}}$ the space formed from $\widetilde{X}$ by coning off hyperplane-carriers; here we imitate that proof. Define a map $f : \Gamma^{(0)} \to \widehat{\mathcal{H}}^{(0)}$ as follows: if v is a vertex of Γ, corresponding to a subcomplex $\widetilde{C}$ (either a lift of some $\widetilde{Y}_i$ or a hyperplane-carrier), then let $f(v) \in \widehat{\mathcal{H}}^{(0)}$ be the cone-point over the corresponding subcomplex. Since each point of $\widetilde{X}$ lies in a hyperplane-carrier, the map f is quasi-surjective, i.e. $\widehat{\mathcal{H}}^{(0)}$ lies in a uniform neighborhood of the image of f. To see that f is a quasi-isometric embedding, one argues exactly as in [Hag14, Section 5]. □

Lemma 5.3 (Loxodromics persist I)**.** *Let $\tilde{g} \in \pi_1 X$ act loxodromically on $\widetilde{\mathcal{H}}$. Then either $\tilde{g}$ acts loxodromically on $\widehat{\mathcal{H}}$ or $\tilde{g}$ stabilizes some elevation $\widetilde{Y}_i \subset \widetilde{X}$ of some $Y_i \to X$.*

Proof. By Lemma 5.2, $\widehat{\mathcal{H}}$ is quasi-isometric to the intersection graph Γ of the set of hyperplane-carriers and lifts of the various $\widetilde{Y}_i$ in $\widetilde{X}$, and this graph is connected. Hence it suffices to show that $\langle\tilde{g}\rangle$ acts loxodromically on Γ provided $\tilde{g} \notin K$.

Let $\widetilde{A} \subseteq \widetilde{X}$ be a combinatorial geodesic axis for $\tilde{g}$ (by replacing $\widetilde{X}$ by its first cubical subdivision, we may assume that such an axis exists [Hag]). Fix a 0–cube $\tilde{a} \in \widetilde{A}$ and fix $n > 0$. Let P be the subpath of $\widetilde{A}$ joining $\tilde{a}$ to $g^n\tilde{a}$.

Let $\widehat{Q}$ be a geodesic of Γ joining vertices corresponding to subcomplexes containing $\tilde{a}$ and $g^n\tilde{a}$. Let the vertex-sequence of $\widehat{Q}$ be $C_0, \dots, C_N$, where each C_i is either a hyperplane-carrier in $\widetilde{X}$ or a lift of some $\widetilde{Y}_i$ to $\widetilde{X}$. Then we have a combinatorial path $Q = \alpha_0 \cdots \alpha_N$ joining $\tilde{a}$ to $g^n\tilde{a}$, where each α_i is a geodesic in C_i. The closed path QP^{-1} bounds a disc diagram $D \to \widetilde{X}$.

Suppose that we have chosen $\widehat{Q}$, and then chosen Q, and then chosen D, so that the area of D is minimal among all possible such choices. Let K be a dual curve in D emanating from P.

Then K ends on Q, since P is a geodesic of $\widetilde{X}$ and combinatorial geodesics are characterized by the fact that they don't contain two 1–cubes dual to the same hyperplane. If K is a dual curve emanating from Q, then it emanates from some α_i. Suppose that K ends on α_j. Then $j \neq i$ since α_i is a geodesic. If $|i-j|>2$, then since the hyperplane to which K maps has carrier Γ–adjacent to C_i, C_j, we have contradicted that $\widehat{Q}$ is a Γ–geodesic. If $j = i \pm 2$, then we can replace $C_{i\pm 1}$ with the carrier of the hyperplane to which K maps, providing a new choice of $\widehat{Q}$ leading to a lower-area choice of D. Finally, if $j = i \pm 1$, then we can apply hexagon moves to show that α_i has a terminal segment coinciding with an initial segment of α_{i+1} (say); we can remove the resulting spur. We conclude that D can be chosen so that all dual curves travel from Q to P. In other words, $\widehat{Q}$ is a geodesic of $\widetilde{X}$.

Since $\tilde{g}$ acts loxodromically on $\widetilde{\mathcal{H}}$, Lemma 5.2 shows that there is a constant $\kappa_0 \geqslant 1$ so that for all hyperplanes H of $\widetilde{X}$, there are at most κ_0 hyperplanes that intersect $\widetilde{A}$ and also intersect H (i.e. $\widetilde{A}$ has uniformly bounded projection to each hyperplane). Hence, for each i for which C_i is a hyperplane-carrier, at most κ_0 dual curves in D travel from P to α_i.

We now claim that there exists κ_1 so that, for all $i \in \mathcal{I}$ and all elevations $\widetilde{Y}_i \subseteq \widetilde{X}$ of $Y_i \to X$, there are at most κ_1 hyperplanes intersecting $\widetilde{Y}_i$ and $\widetilde{A}$. First, since g is loxodromic on $\widetilde{\mathcal{H}}$, Lemma 5.2 and Theorem 2.3 of [Hag13] imply that there exists κ' such that if $\mathsf{d}_{\widetilde{X}}(H\cap\widetilde{A}, H'\cap\widetilde{A}) > \kappa'$, for hyperplanes H, H', then $H \cap H' = \emptyset$. It follows that the cubical convex hull $\widetilde{B}$ of $\widetilde{A}$ lies at finite Hausdorff distance from $\widetilde{A}$. Suppose that no κ_1 with the desired property exists. Then for any $L > 0$, there exists $\widetilde{Y}_i$ so that $\operatorname{diam}(\mathfrak{g}_{\widetilde{B}}(\widetilde{Y}_i)) > L$, where $\mathfrak{g}_{\widetilde{B}} : \widetilde{X} \to \widetilde{B}$ is the combinatorial closest-point projection, or *gate* map; see e.g. [BHS14]. Translating by elements of $\langle \tilde{g} \rangle$ and applying the $C'(\frac{1}{144})$ condition now shows that there exists a unique $\widetilde{Y}_i$ so that $\widetilde{A}$ lies in $\mathcal{N}_r(\widetilde{Y}_i)$ for some $r \geqslant 0$. Since the same is true of $\tilde{g}\widetilde{A} = \widetilde{A}$ and $\tilde{g}\widetilde{Y}_i$, we have $\tilde{g} \in \operatorname{Stab}(\widetilde{Y}_i)$.

We have shown that if $\tilde{g}$ does not stabilize some elevation of some Y_i, then $\mathsf{d}_{\widetilde{X}}(\tilde{a}, g^n\tilde{a}) \leqslant (\kappa_0+\kappa_1)N \leqslant (\kappa_0+\kappa_1)(\lambda \mathsf{d}_{\widehat{\mathcal{H}}}(\tilde{a}, g^n\tilde{a}) + \mu)$, where λ, μ are constants for the quasi-isometry $\Gamma \to \widehat{\mathcal{H}}$. Since the left-hand side grows linearly in n, it follows that $\tilde{g}$ is loxodromic on $\widehat{\mathcal{H}}$. □

Lemma 5.4 (Loxodromics persist II)**.** *Let $\tilde{g} \in \pi_1 X$ act loxodromically on $\widetilde{\mathcal{H}}$ and let $g = \tilde{g}K \in \pi_1 X^*$. Then either g acts loxodromically on $\mathcal{H}$ or $\tilde{g}$ stabilizes an elevation to $\widetilde{X}$ of some $Y_i \to X$.*

Proof. Each lift $Y_i \to \widetilde{X}^*$ is an embedding, by the small-cancellation conditions and [Wis, Theorem 4.1]. It suffices to show that if g is not loxodromic on $\mathcal{H}$, then g stabilizes some lift to $\widetilde{X}^*$ of some Y_i. Note that for each i, the elevations $\widetilde{Y}_i$ to $\widetilde{X}$ of $Y_i \to X$ correspond to the lifts to $\widetilde{X}^*$ of $Y_i \to X$.

Let $\widetilde{A}$ be a combinatorial axis for $\tilde{g}$ in $\widetilde{X}$ and let $\tilde{a} \in \widetilde{A}$ be a 0–cube. Recall that $p : \widetilde{X} \to \operatorname{Cay}(X^*)$ is the universal covering map. Let $a = p(\tilde{a})$, so that $g^n a = p(\tilde{g}^n\tilde{a})$ for each $n > 0$. Fixing $n > 0$, choose a sequence $C_0, \dots, C_N$ satisfying:

- for each $j \leqslant N$, either C_j is the carrier of a hyperplane in $\widetilde{X}^*$ or $C_j \to \widetilde{X}^*$ is a lift of some $Y_i \to X$ (abusing notation, we use the same name for C_j as for its image);
- $a \in C_0$ and $g^n a \in C_N$;
- $C_j \cap C_{j+1} \neq \emptyset$ for all $j \leqslant N$;
- N is minimal with the above properties.

Let P be a combinatorial geodesic of $\operatorname{Cay}(X^*)$ joining a to $g^n a$. For each j, let α_i be a combinatorial geodesic of C_j so that $Q = \alpha_0\alpha_2\cdots\alpha_N$ is a piecewise-geodesic from a to $g^n a$ with no backtracks. Let $D \to \widetilde{X}^*$ be a minimal disc diagram with boundary path PQ^{-1}. Moreover, choose the C_j, P, and Q subject to the above constraints so that D has minimal area among all diagrams constructed in the preceding manner.

By Theorem 2.15 and the fact that P is a geodesic, D does not contain a positively-curved shell whose outer path is a subpath of P. Likewise, suppose that K is a positively-curved shell

with boundary path OI, where the out path O is a subpath of Q. Then the relator Y to which K maps intersects C_i, C_j for some i, j, and we can replace O by I in Q to obtain a new collection of data of the above type leading to a disc diagram which is a subdiagram of D with one fewer cone-cells, contradicting minimality. Hence, if there is a positively-curved shell in D, its outer path lies at the transition from P to Q. The same holds for spurs by construction.

By performing square homotopies, we see that D decomposes as $D = D_0 \cup_{P'} L \cup_{P''} D_1$, where $D_0, D_1 \to \widetilde{X}^*$ are square diagrams (i.e. factor as $D_0, D_1 \to \widetilde{X} \to \widetilde{X}^*$) bounded by geodesic bigons $\partial_p D_0 = PP'$ and $\partial_p D_1 = P''Q$, and no generalized corners in L lie along P', P''. The subdiagram L is bounded by $P'P''$. Theorem 2.13 shows that L is a ladder; let $\{R_k\}$ and $\{K_k\}$ respectively denote the constituent pseudorectangles and cone-cells of L. See Figure 9.

Since $\tilde{g}$ is rank-one (because it is loxodromic on $\widetilde{\mathcal{H}}$, while half-flats project to diameter–$\leqslant 3$ sets in the contact graph), and since D_0, D_1 lift to $\widetilde{X}$, there exists κ_0, independent of n, so that P, P' lie at Hausdorff distance at most κ_0, and the same is true of P'', Q. In other words, the square bigons D_0, D_1 are κ_0–thin.

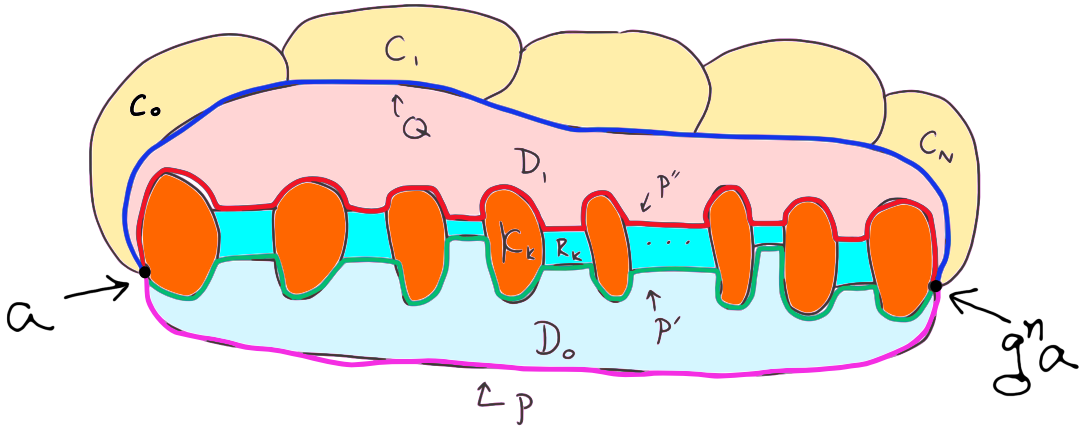

FIGURE 9. The diagram D and its environs, in the proof of Lemma 5.4.

Let $r = |\{R_k\}|$ denote the number of constituent pseudorectangles of L. Then there exists K_1, independent of n, so that $r \geqslant K_1 n$. Indeed, this can be seen by considering a lift of D_0 or D_1 to $\widetilde{X}$ and applying Lemma 5.3.

Lemma 3.1 and the fact that g does not stabilize a lift of a relator combine to show that there is a uniform bound (independent of n) on the number of pseudorectangles of L whose boundary paths can come κ_0–close to a given C_j. Otherwise, we would find (analogously to the proof of Lemma 5.3) some $Y_i \to \mathrm{Cay}(X^*)$ with $\langle g \rangle a$ lying in a fixed neighborhood of Y_i. The small-cancellation assumptions would then show that $gY_i = Y_i$.

Combining the above two conclusions shows that N grows linearly in n, whence g is loxodromic on $\mathcal{H}$, as required, unless g stabilizes some Y_i. □

5.2. **Building loxodromics in practice.** In this section, we discuss how to find $\tilde{g} \in \pi_1 X$ so that $\tilde{g}$ acts loxodromically on $\widetilde{\mathcal{H}}$ without stabilizing any of the $\widetilde{Y}_i$.

We now introduce the assumption that X is compact (although we still put no restriction on the set of relators $Y_i \to X$). We also assume, for the moment, that $\pi_1 X$ acts *essentially* on $\widetilde{X}$, in the sense that $\widetilde{X}$ does not contain a $\pi_1 X$–invariant proper convex subcomplex. This is not a real restriction: in our application, we will be able to reduce to this case. These assumptions are imposed in order to use the *double-skewering lemma* of Caprace-Sageev [CS11], of which we use the following form from [Hag13]:

Lemma 5.5 (Strong double-skewering)**.** *Let G act properly, essentially, and cocompactly on a CAT(0) cube complex $\widetilde{X}$. Suppose H, V are hyperplanes with $\mathsf{d}_{\mathcal{C}\widetilde{X}}(H, V) \geqslant 3$. Then there exists $g \in G$ so that V separates H, gH, and g acts loxodromically on $\mathcal{C}\widetilde{X}$.*

Let $\widetilde{Y}_i \hookrightarrow \widetilde{X}$ be an elevation of some $Y_i \to X$, and let $\mathcal{C}\widetilde{Y}_i$ be the subgraph of the contact graph $\mathcal{C} = \mathcal{C}\widetilde{X}$ of $\widetilde{X}$ spanned by the vertices corresponding to the hyperplanes intersecting the convex subcomplex $\widetilde{Y}_i$ (so $\mathcal{C}\widetilde{Y}_i$ is a copy of the contact graph of $\widetilde{Y}_i$). For convenience, we now denote by $\mathcal{R}$ the collection of all subcomplexes $\widetilde{Y}_i$ of $\widetilde{X}$ which are elevations of relators $Y_i \to X$.

The following result reflects a general small-cancellation intuition. It is of independent interest and might be useful even in classical and certain graphical small-cancellation presentations. We denote by $\mathcal{N}_r(\mathcal{C}\widetilde{Y})$ the closed r-neighborhood of $\mathcal{C}\widetilde{Y}$.

Lemma 5.6 (Quasiconvex malnormal collection)**.** *Let $\widetilde{S}$ denote the set of all elevations $\widetilde{Y}_i$ to $\widetilde{X}$ of relators. Then there exists $\Delta : [0, \infty) \to [0, \infty)$ so that:*

(1) $\mathcal{C}\widetilde{Y}$ is $\Delta(0)$–quasiconvex in $\mathcal{C}$ for all $\widetilde{Y} \in \mathcal{R}$.
(2) If $\widetilde{Y}, \widetilde{Y}' \in \mathcal{R}$ are distinct, then for all $r \geqslant 0$,

$$\operatorname{diam}(\mathcal{N}_r(\mathcal{C}\widetilde{Y}) \cap \mathcal{C}\widetilde{Y}') \leqslant \Delta(r).$$

Remark 5.7. The above lemma does not use compactness of X.

Proof of Lemma 5.6. We first prove assertion (1). Let $H, V \in \mathcal{C}\widetilde{Y}$ be vertices (i.e. hyperplanes crossing $\widetilde{Y}$) and let $x, y \in \widetilde{Y}$ be 0–cubes in the carriers of H, V respectively. Then by Proposition 3.1 of [BHS14], there exists a geodesic $H = H_0, H_1, \dots, H_n = V$ of $\mathcal{C}$ and a geodesic $\gamma = \gamma_0\gamma_1 \cdots \gamma_n$ of $\widetilde{X}$ that joins x to y and has the property that γ_i lies in the carrier of H_i for each i. Now, since $\widetilde{Y}$ is convex in $\widetilde{X}$, each hyperplane crossing γ corresponds to a vertex of $\mathcal{C}\widetilde{Y}$, and hence $\mathsf{d}_{\mathcal{C}}(H_i, \mathcal{C}\widetilde{Y}_i) \leqslant 1$ for each i, i.e. $\mathcal{C}$ contains a geodesic joining H, V and lying 1–close to $\mathcal{C}\widetilde{Y}$; the claim now follows since $\mathcal{C}$ is hyperbolic.

We now prove (2). Let $\widetilde{Y}, \widetilde{Y}'$ be distinct and fix $r \geqslant 0$. Choose hyperplanes H', V' that cross $\widetilde{Y}'$ and correspond to vertices in $\mathcal{C}\widetilde{Y}'$ whose closest vertices in $\mathcal{C}Y$, denoted H, V respectively, satisfy $\mathsf{d}_{\mathcal{C}}(H, H') \leqslant r, \mathsf{d}_{\mathcal{C}}(V, V') \leqslant r$.

Choose $\mathcal{C}$–geodesic sequences of hyperplanes $H = H_0, \dots, H_s = H', V = V_0, \dots, V_t = V'$ with $s, t \leqslant r$. We must show that $\mathsf{d}_{\mathcal{C}}(H', V')$ is bounded in terms of r. At the expense of adding $2r$ to our eventual bound, we may assume that none of the H_i or V_j, with $i < s, j < t$, crosses $\widetilde{Y}'$.

Let $P, P' \to \widetilde{Y}, \widetilde{Y}'$ be combinatorial geodesics joining $\mathcal{N}(V), \mathcal{N}(H)$ and $\mathcal{N}(H'), \mathcal{N}(V')$ respectively. For each i, let $Q_i \to \mathcal{N}(H_i)$ be a geodesic, and $R_i \to \mathcal{N}(V_i)$ be a geodesic, so that $S = P \prod_i Q_i P' \prod_j R_j$ is a closed path in $\widetilde{X}$. Let $D \to \widetilde{X}$ be a minimal-area disc diagram bounded by S, and suppose that all of the above paths and hyperplanes were chosen (fixing H, H', V, V') so as to minimize the area of the resulting diagram D.

Note that $\mathsf{d}_{\mathcal{C}}(H', V')$ is bounded by the number of dual curves in D emanating from P'. Let K be such a dual curve. By minimality of the area of D (allowing P' and the Q_i, R_i to vary), we can assume that K does not end on Q_s or R_t. Hence K ends either on P, on Q_i for $0 \leqslant i < s$, or on R_j for $0 \leqslant j < t$. Let $M \geqslant 0$ bound the lengths of geodesic cone-pieces and wall-pieces; such a bound exists by our small-cancellation assumptions. Then the number of dual curves traveling from P' to P is at most M, since P, P' are geodesics and thus any two such curves map to distinct hyperplanes. Next, since H_i (respectively V_j) has the property that H_i (respectively V_j) does not cross $\widetilde{Y}'$, then there are at most M dual curves traveling from P' to Q_i (respectively R_j). This completes the proof, with $\Delta(r) \leqslant 2Mr + 2r$. □

Lemma 5.8. *There exists $R \geqslant 0$ so that the following holds. Suppose that there exist distinct $\widetilde{Y}, \widetilde{Y}' \in \mathcal{R}$ with $\mathcal{C}\widetilde{Y}, \mathcal{C}\widetilde{Y}'$ subgraphs of $\mathcal{C}$ of diameter at least $10R$. Then there exists $\tilde{g} \in \pi_1 X$ acting loxodromically on $\mathcal{C}$ and not stabilizing any element of $\mathcal{R}$.*

Proof. Let $\pi\colon \mathcal{C} \to \mathcal{C}\widetilde{Y}$ be the coarse closest-point projection, which exists since $\mathcal{C}\widetilde{Y}$ is uniformly quasiconvex by Lemma 5.6. Let π' be the corresponding projection to $\mathcal{C}\widetilde{Y}'$. Lemma 5.6 implies that $\pi(\mathcal{C}\widetilde{Y}')$ has diameter bounded by some Δ depending on the hyperbolicity constant of $\mathcal{C}$ and the quasiconvexity constant $\Delta(0)$, and the same is true of $\pi'(\mathcal{C}\widetilde{Y})$.

Let R be a constant to be determined, and choose vertices $H, H' \in \mathcal{C}\widetilde{Y}, \mathcal{C}\widetilde{Y}'$, respectively, so that $\mathsf{d}_{\mathcal{C}}(H, \pi(\mathcal{C}\widetilde{Y}')) \geqslant R$ and $\mathsf{d}_{\mathcal{C}}(H', \pi'(\mathcal{C}(\widetilde{Y}))) \geqslant R$. Let $\gamma \to \mathcal{C}$ be a geodesic joining H, H'. Then thin quadrilateral considerations show that γ passes uniformly close to $\pi(\mathcal{C}\widetilde{Y}')$ and to $\pi'(\mathcal{C}\widetilde{Y})$, so that $\mathsf{d}_{\mathcal{C}}(H, H') > 2$ provided R is sufficiently large.

Now, since $\mathsf{d}_{\mathcal{C}}(H, H') > 2$, the hyperplanes H, H' of $\widetilde{X}$ satisfy the hypotheses of Lemma 5.5, so that $\pi_1 X$ contains an element $\tilde{g}$ whose axis $\widetilde{A}$ in $\widetilde{X}$ passes through both hyperplanes H, H', and projects to a quasigeodesic axis $\overline{A}$ in $\mathcal{C}$ containing the vertices H, H'. Hence either $\tilde{g}$ has the desired property, or $\tilde{g}$ stabilizes some $\widetilde{Y}'' \in \mathcal{R} - \{\widetilde{Y}, \widetilde{Y}'\}$.

In the latter case, every hyperplane crossing $\widetilde{A}$ crosses $\widetilde{Y}''$. Let α be the subpath of $\widetilde{A}$ between the 1–cubes dual to H, H', let the geodesic $\beta \to \mathcal{N}(H')$ join the endpoint of α to a closest vertex of $\mathcal{N}(H') \cap \widetilde{Y}'$, let the geodesic $\gamma \to \widetilde{Y}'$ join the endpoint of β to a point $y' \in \widetilde{Y}'$ as close as possible to $\widetilde{Y}$, let the geodesic ζ join y' to a closest point $y \in \widetilde{Y}$, let $\eta \to \widetilde{Y}'$ join y to a closest point of $\mathcal{N}(H)$, let $\theta \to \mathcal{N}(H)$ join the endpoint of η to the initial point of α, and let $D \to \widetilde{X}$ be a disc diagram with boundary path $\alpha\beta\gamma\zeta\eta\theta$, with minimal area for its boundary path and with each of the named geodesics chosen in its path-homotopy class so that the resulting diagram is minimal. Observe that $|\eta|, |\gamma| \geqslant R$ since these paths must project to paths in $\mathcal{C}$ of length at least R. Now, no dual curve in D travels from $\eta \cup \gamma$ to ζ because such dual curves map to hyperplanes crossing $\widetilde{Y} \cup \widetilde{Y}'$, while the hyperplanes crossing ζ are exactly those separating $\widetilde{Y}, \widetilde{Y}'$, by our choice of y, y'. Similarly, no dual curve travels from η to θ or γ to β. Hence at least $\frac{R}{3}$ dual curves emanating from η travel to α, or to γ, or to β. In either of the former cases, $\widetilde{Y}$ contains a cone-piece of length $\frac{R}{3}$, provided $R > 3M$, where M is the bound on geodesic pieces coming from the small-cancellation conditions. Otherwise, since H' cannot cross $\widetilde{Y}$ – it was chosen to be very $\mathcal{C}$–distant from $\mathcal{C}\widetilde{Y}$ – β provides a wall-piece in $\widetilde{Y}$ of length $\frac{R}{3}$, which is again impossible if $R > 3M$. This completes the proof. □

Lemma 5.9. *Let $\tilde{g}$ act loxodromically on $\mathcal{C}$ and suppose that there exists B such that $\operatorname{diam}(\mathcal{C}\widetilde{Y}) \leqslant B$ for all $\widetilde{Y} \in \mathcal{R}$. Then $\tilde{g}$ does not stabilize any element of $\mathcal{R}$.*

Proof. This is obvious. □

Definition 5.10 (Normalized)**.** The cubical presentation $\langle X \mid \{Y_i\}_{i\in\mathcal{I}}\rangle$ is *normalized* if the nonpositively-curved cube complex Y_i is not contractible for each $i \in \mathcal{I}$.

Lemma 5.11 (Building loxodromics)**.** *Suppose that $\langle X \mid \{Y_i\}_{i\in\mathcal{I}}\rangle$ is a normalized $C'(\frac{1}{144})$ cubical presentation with X compact and each Y_i compact. Then one of the following holds:*

(A) $\pi_1 X^*$ is finite;
(B) $\mathcal{I} = \emptyset$, and $\pi_1 X^* = \pi_1 X$, and $\widetilde{X}$ contains a $\pi_1 X$–invariant convex subcomplex decomposing as the product of two unbounded subcomplexes;
(C) there exists $\tilde{g} \in \pi_1 X$ acting loxodromically on $\widetilde{\mathcal{H}}$ and not stabilizing any $\widetilde{Y}_i$.

Proof. We first reduce to the case where $\pi_1 X$ acts essentially on $\widetilde{X}$, i.e. there is no convex, $\pi_1 X$–invariant proper subcomplex in $\widetilde{X}$. Indeed, by Proposition 3.5 of [CS11], there is a $\pi_1 X$–invariant, $\pi_1 X$–essential convex subcomplex $\widetilde{Z} \subseteq \widetilde{X}$. Let Z denote the quotient of $\widetilde{Z}$ by $\pi_1 X$, so

that $\widetilde{Z} \hookrightarrow \widetilde{X}$ descends to a local isometry $Z \to X$. We form a new cubical presentation for $\pi_1 X^*$ by attaching to Z all of the components of each fiber product $Y_i \otimes_X Z \to Z$ and then normalizing (discarding contractible relators). Pieces in the new cubical presentation arise by intersecting pieces in $\widetilde{X}$ with $\widetilde{Z}$, so $C'(\frac{1}{144})$ condition persists. Thus, assume $\pi_1 X$ acts essentially on $\widetilde{X}$.

If $\mathcal{I} = \emptyset$, then either (B) holds, or the Irreducibility Criterion of [CS11] provides hyperplanes H, H' at $\mathcal{C}$–distance at least 3, whence Lemma 5.5 and Lemma 5.2 show that (C) holds.

If there exist distinct $\widetilde{Y}, \widetilde{Y}' \in \mathcal{R}$ with $\mathcal{C}\widetilde{Y}, \mathcal{C}\widetilde{Y}'$ subgraphs of $\mathcal{C}$ of diameter at least $10R$, then Lemma 5.8 provides $\tilde{g} \in \pi_1 X$ so that $\tilde{g}$ does not stabilize an element of $\mathcal{R}$ and acts loxodromically on $\mathcal{C}$, and therefore, by Lemma 5.2, on $\widetilde{\mathcal{H}}$. Thus conclusion (C) holds in this case. If $\widetilde{Y} \in \mathcal{R}$ has the property that $\mathcal{C}\widetilde{Y}$ has diameter $\geqslant 10R$, and $\widetilde{Y}$ is unique with this property, then $g\widetilde{Y} = \widetilde{Y}$ for all $g \in \pi_1 X$, and conclusion (A) holds. Indeed, since Y is compact and $\widetilde{Y}$ is $\pi_1 X$–invariant, $\pi_1 Y$ has finite index in $\pi_1 X$. (In fact, essentiality shows $\widetilde{X} = \widetilde{Y}$.)

It remains to consider the case in which $\mathcal{C}\widetilde{Y}$ is uniformly bounded for all $\widetilde{Y}$. Lemma 5.9 and Lemma 5.2 show that conclusion (C) holds exactly if there exists $\tilde{g}$ acting loxodromically on $\mathcal{C}$.

Applying the Irreducibility Criterion and Lemma 5.5 as above show that if conclusion (C) does not hold, then $\widetilde{X}$ decomposes as a product $A \times B$ of unbounded subcomplexes. In particular, hyperplanes have the form $A \times H$ or $V \times B$, where H, V are respectively hyperplanes of B, A. In this case, we claim that either conclusion (B) or (A) holds. Indeed, suppose $\widetilde{Y} \in \mathcal{R}$. Then $\widetilde{Y} = A' \times B'$, where $A', B' \subset A, B$ are convex subcomplexes.

If $A' = A, B' = B$, then (A) holds. Hence let $V \times B$ be a hyperplane of $\widetilde{X}$ that does not cross $\widetilde{Y}$, where V is a hyperplane of A not crossing A'. Then any geodesic in B' is a wall-piece, so $\operatorname{diam}(B') \leqslant M$. If there is also a hyperplane $A \times H$ so that H does not cross B', then we can conclude that $\operatorname{diam}(A') \leqslant M$, whence $\operatorname{diam}(\widetilde{Y}) \leqslant 2M$. In particular, since we can take M to be $\leqslant \frac{1}{144}^{th}$ of the length of any geodesic in $\widetilde{Y}$ projecting to an essential closed path in X, there is no essential closed path in Y, so Y is contractible; since our cubical presentation was normalized, Y cannot exist. Otherwise, $\operatorname{diam}(B') \leqslant M$ and $A = A'$, which means that there are infinite wall-pieces, and infinite cone-pieces between $\widetilde{Y}$ and its translates, so such $\widetilde{Y}$ cannot exist. Thus (B) holds. This completes the proof. □

5.3. **The proof.** We conclude with:

Proof of Theorem B. We are interested in a statement about $\pi_1 X^*$, and discarding relators does not weaken small-cancellation conditions, so we assume that $\langle X \mid \{Y_i\}_{i \in \mathcal{I}} \rangle$ is normalized. The theorem now follows from Lemma 5.11, Lemma 5.4, and Theorem 4.3. □

References

[AD12] Goulnara Arzhantseva and Cornelia Drutu. Geometry of infinitely presented small cancellation groups, Rapid Decay and quasi-homomorphisms. arXiv:1212.5280, 2012.

[AMS] Yago Antolin, Ashot Minasyan, and Alessandro Sisto. Commensurating endomorphisms of acylindrically hyperbolic groups and applications. *Groups, Geometry, and Dynamics, to appear.* arXiv:1310.8605, 2013.

[AO14] Goulnara Arzhantseva and Damian Osajda. Graphical small cancellation groups with the Haagerup property. arXiv:1404.6807, 2014.

[AS83] Kenneth I. Appel and Paul E. Schupp. Artin groups and infinite Coxeter groups. *Invent. Math.*, 72(2):201–220, 1983.

[Bal16] Sahana Balasubramanya. Acylindrical group actions on quasi-trees. arXiv:1602.03941, 2016.

[BB96] Werner Ballmann and Serguei Buyalo. Nonpositively curved metrics on 2-polyhedra. *Mathematische Zeitschrift*, 222(1):97–134, 1996.

[BF02] Mladen Bestvina and Koji Fujiwara. Bounded cohomology of subgroups of mapping class groups. *Geometry & Topology*, 6(1):69–89, 2002.

[BF08] Mladen Bestvina and Mark Feighn. A hyperbolic $\mathrm{Out}(F_n)$-complex. arXiv:0808.3730, 2008.

[BHS14] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces I: curve complexes for cubical groups. arXiv:1412.2171, 2014.

[Bow08] Brian H Bowditch. Tight geodesics in the curve complex. *Inventiones mathematicae*, 171(2):281–300, 2008.

[Bow14] Brian H Bowditch. Uniform hyperbolicity of the curve graphs. *Pacific J. Math*, 269(2):269–280, 2014.

[Bri48] I.A. Brin. Gauss-Bonnet theorems for polyhedra. *Uspekhi Mat. Nauk*, 3:226–227, 1948.

[Che00] Victor Chepoi. Graphs of some CAT(0) complexes. *Advances in Applied Mathematics*, 24(2):125–179, 2000.

[CS11] Pierre-Emmanuel Caprace and Michah Sageev. Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.*, 21:851–891, 2011.

[DGO11] François Dahmani, Vincent Guirardel, and Denis Osin. Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. arXiv:1111.7048, 2011.

[FPS15] Roberto Frigerio, Beatrice M Pozzetti, and Alessandro Sisto. Extending higher-dimensional quasicocycles. *Journal of Topology*, 8(4):1123–1155, 2015.

[Ger87] Steve M Gersten. *Reducible diagrams and equations over groups*. Springer, 1987.

[Gro03] Mikhail Gromov. Random walk in random groups. *Geometric and Functional Analysis*, 13(1):73–146, 2003.

[Gru] Dominik Gruber. Infinitely presented $C(6)$-groups are SQ-universal. *J. Lond. Math. Soc. (2)*, 92(1):178–201.

[Gru15] Dominik Gruber. *Infinitely presented graphical small-cancellation groups*. PhD thesis, Universität Wien, 2015.

[GS14] Dominik Gruber and Alessandro Sisto. Infinitely presented graphical small cancellation groups are acylindrically hyperbolic. arXiv:1408.4488, 2014.

[Hag] Frédéric Haglund. Isometries of CAT(0) cube complexes are semi-simple. arXiv:0705.3386, 2007.

[Hag13] Mark F Hagen. The simplicial boundary of a CAT(0) cube complex. *Alg. Geom. Topol.*, 13:1299–1367, 2013.

[Hag14] Mark F Hagen. Weak hyperbolicity of cube complexes and quasi-arboreal groups. *J. Topol.*, 7(2):385–418, 2014.

[Ham07] Ursula Hamenstadt. Geometry of the complex of curves and of Teichmüller space. *Handbook of Teichmüller theory*, 1:447–467, 2007.

[Ham08] Ursula Hamenstädt. Bounded cohomology and isometry groups of hyperbolic spaces. *Journal of the European Mathematical Society*, 10(2):315–349, 2008.

[HO13] Michael Hull and Denis Osin. Induced quasicocycles on groups with hyperbolically embedded subgroups. *Algebraic & Geometric Topology*, 13(5):2635–2665, 2013.

[HW08] Frédéric Haglund and Daniel T Wise. Special cube complexes. *Geometric and Functional Analysis*, 17(5):1551–1620, 2008.

[Jan14] Kasia Jankiewicz. Greendlinger's lemma in cubical small-cancellation theory. arXiv:1401.4995, 2014.

[Mac] John M. Mackay. Conformal dimension and random groups. *Geom. Funct. Anal.*, 22(1):213–239.

[MO15] Ashot Minasyan and Denis Osin. Acylindrical hyperbolicity of groups acting on trees. *Mathematische Annalen*, 362(3-4):1055–1105, 2015.

[MW02] Jonathan P McCammond and Daniel T Wise. Fans and ladders in small cancellation theory. *Proceedings of the London Mathematical Society*, 84(3):599–644, 2002.

[NR98] Graham Niblo and Martin Roller. Groups acting on cubes and Kazhdan's property (T). *Proceedings of the American Mathematical Society*, 126(3):693–699, 1998.

[Oll06] Yann Ollivier. On a small cancellation theorem of Gromov. *Bulletin of the Belgian Mathematical Society Simon Stevin*, 13(1):75–90, 2006.

[Osi15a] Denis Osin. Acylindrically hyperbolic groups. *Transactions of the American Mathematical Society*, 2015.

[Osi15b] Denis Osin. On acylindrical hyperbolicity of groups with positive first ℓ^2-Betti number. arXiv:1501.03066, 2015.

[OW07] Yann Ollivier and Daniel Wise. Kazhdan groups with infinite outer automorphism group. *Transactions of the American Mathematical Society*, 359(5):1959–1976, 2007.

[PS15] Piotr Przytycki and Alessandro Sisto. A note on acylindrical hyperbolicity of mapping class groups. arXiv:1502.02176, 2015.

[RS87] Eliyahu Rips and Yoav Segev. Torsion-free group without unique product property. *Journal of Algebra*, 108(1):116–126, 1987.

[Sam] Andrea Sambusetti. Growth tightness of surface groups. *Expo. Math.*, 20(4):345–363.

[Sis] Alessandro Sisto. Contracting elements and random walks. *Journal für die reine und angewandte Mathematik, to appear*. arXiv:1112.2666, 2011.

[Sis14] Alessandro Sisto. Quasi-convexity of hyperbolically embedded subgroups. *Mathematische Zeitschrift*, pages 1–10, 2014.

[Str90] Ralph Strebel. Appendix. Small cancellation groups. In *Sur les groupes hyperboliques d'après Mikhael Gromov (Bern, 1988)*, volume 83 of *Progr. Math.*, pages 227–273. Birkhäuser Boston, Boston, MA, 1990.

[Wis] Daniel T. Wise. The structure of groups with a quasiconvex hierarchy. 205 pp., preprint 2011.

[Wis04] Daniel T Wise. Cubulating small cancellation groups. *Geometric & Functional Analysis*, 14(1):150–214, 2004.

[Wis12] Daniel Wise. *From Riches to Raags: 3-manifolds, right-angled Artin Groups, and Cubical Geometry*, volume 117. American Mathematical Society, 2012.

University of Vienna, Faculty of Mathematics, Oskar-Morgenstern-Platz 1, 1090 Wien, Austria

E-mail address: goulnara.arzhantseva@univie.ac.at

Department of Pure Maths. and Math. Stats., University of Cambridge, Cambridge, UK

E-mail address: markfhagen@gmail.com